Exam Questions DVA-C02

DVA-C02

https://www.2passeasy.com/dumps/DVA-C02/

NEW QUESTION 1
A developer wants to add request validation to a production environment Amazon API Gateway API. The developer needs to test the changes before the API is deployed to the production environment. For the test, the developer will send test requests to the API through a testing tool.
Which solution will meet these requirements with the LEAST operational overhead?

A. Export the existing API to an OpenAPI fil
B. Create a new AP
C. Import the OpenAPI file.
D. Perform the test   Modify the new API to add request validatio
E. Modify the existing API to add request validatio
F. Deploy the existing API to production.
G. Modify the existing API to add request validatio
H. Deploy the updated API to a new API Gateway stag
I. Perform the test
J. Deploy the updated API to the API Gateway production stage.
K. Create a new AP
L. Add the necessary resources and methods, including new request validatio
M. Perform the test
N. Modify the existing API to add request validatio
O. Deploy the existing API to production.
P. Clone the existing AP
Q. Modify the new API to add request validatio
R. Perform the test
S. Modify the existing API to add request validatio
T. Deploy the existing API to production.

Answer: B

Explanation:
Amazon API Gateway allows you to create, deploy, and manage a RESTful API to expose backend HTTP endpoints, AWS Lambda functions, or other AWS services1. You can use API Gateway to perform basic validation of an API request before proceeding with the integration request1. When the validation fails, API Gateway immediately fails the request, returns a 400 error response to the caller, and publishes the validation results in CloudWatch Logs1.
To test changes before deploying to a production environment, you can modify the existing API to add request validation and deploy the updated API to a new API Gateway stage1. This allows you to perform tests without affecting the production environment. Once testing is complete and successful, you can then deploy the updated API to the API Gateway production stage1.
This approach has the least operational overhead as it avoids unnecessary creation of new APIs or exporting and importing of APIs. It leverages the existing infrastructure and only requires changes in the configuration of the existing API1.

NEW QUESTION 2
A developer is creating a mobile app that calls a backend service by using an Amazon API Gateway REST API. For integration testing during the development phase, the developer wants to simulate different backend responses without invoking the backend service.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AWS Lambda functio
B. Use API Gateway proxy integration to return constant HTTP responses.
C. Create an Amazon EC2 instance that serves the backend REST API by using an AWS CloudFormation template.
D. Customize the API Gateway stage to select a response type based on the request.
E. Use a request mapping template to select the mock integration response.

Answer: D

Explanation:
Amazon API Gateway supports mock integration responses, which are predefined responses that can be returned without sending requests to a backend service. Mock integration responses can be used for testing or prototyping purposes, or for simulating different backend responses based on certain conditions. A request mapping template can be used to select a mock integration response based on an expression that evaluates some aspects of the request, such as headers, query strings, or body content. This solution does not require any additional resources or code changes and has the least operational overhead. Reference: Set up mock integrations for an API Gateway REST API
https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-mock- integration.html

NEW QUESTION 3
A developer is deploying a company's application to Amazon EC2 instances The application generates gigabytes of data files each day The files are rarely accessed but the files must be available to the application's users within minutes of a request during the first year of storage The company must retain the files for 7 years.
How can the developer implement the application to meet these requirements MOST cost- effectively?

A. Store the files in an Amazon S3 bucket Use the S3 Glacier Instant Retrieval storage class Create an S3 Lifecycle policy to transition the files to the S3 Glacier Deep Archive storage class after 1 year
B. Store the files in an Amazon S3 bucke
C. Use the S3 Standard storage clas
D. Create an S3 Lifecycle policy to transition the files to the S3 Glacier Flexible Retrieval storage class after 1 year.
E. Store the files on an Amazon Elastic Block Store (Amazon EBS) volume Use Amazon Data Lifecycle Manager (Amazon DLM) to create snapshots of the EBS volumes and to store those snapshots in Amazon S3
F. Store the files on an Amazon Elastic File System (Amazon EFS) moun
G. Configure EFS lifecycle management to transition the files to the EFS Standard-Infrequent Access (Standard-IA) storage class after 1 year.

Answer: A

Explanation:
Amazon S3 Glacier Instant Retrieval is an archive storage class that delivers the lowest-cost storage for long-lived data that is rarely accessed and requires

retrieval in
milliseconds. With S3 Glacier Instant Retrieval, you can save up to 68% on storage costs compared to using the S3 Standard-Infrequent Access (S3 Standard-IA) storage class, when your data is accessed once per quarter. https://aws.amazon.com/s3/storage- classes/glacier/instant-retrieval/

NEW QUESTION 4
A developer needs to perform geographic load testing of an API. The developer must deploy resources to multiple AWS Regions to support the load testing of the API.
How can the developer meet these requirements without additional application code?

A. Create and deploy an AWS Lambda function in each desired Regio
B. Configure the Lambda function to create a stack from an AWS CloudFormation template in that Region when the function is invoked.
C̣ Create an AWS CloudFormation template that defines the load test resource
D: Use the AWS CLI create-stack-set command to create a stack set in the desired Regions.
E. Create an AWS Systems Manager document that defines the resource
F. Use the document to create the resources in the desired Regions.
G. Create an AWS CloudFormation template that defines the load test resource
H. Use the AWS CLI deploy command to create a stack from the template in each Region.

Answer: B

Explanation:
AWS CloudFormation is a service that allows developers to model and provision AWS resources using templates. A CloudFormation template can define the load test resources, such as EC2 instances, load balancers, and Auto Scaling groups. A CloudFormation stack set is a collection of stacks that can be created and managed from a single template in multiple Regions and accounts. The AWS CLI create-stack-set command can be used to create a stack set from a template and specify the Regions where the stacks should be created. Reference: Working with AWS CloudFormation stack sets

NEW QUESTION 5
A company has an application that uses Amazon Cognito user pools as an identity provider. The company must secure access to user records. The company has set up multi-factor authentication (MFA). The company also wants to send a login activity notification by email every time a user logs in.
What is the MOST operationally efficient solution that meets this requirement?

A. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notificatio
B. Add an Amazon API Gateway API to invoke the functio
C. Call the API from the client side when login confirmation is received.
D. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notificatio
E. Add an Amazon Cognito post authentication Lambda trigger for the function.
F. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notificatio
G. Create an Amazon CloudWatch Logs log subscription filter to invoke the function based on the login status.
H. Configure Amazon Cognito to stream all logs to Amazon Kinesis Data Firehos
I. Create an AWS Lambda function to process the streamed logs and to send the email notification based on the login status of each user.

Answer: B

Explanation:
Amazon Cognito user pools support Lambda triggers, which are custom functions that can be executed at various stages of the user pool workflow. A post authentication Lambda trigger can be used to perform custom actions after a user is authenticated, such as sending an email notification. Amazon SES is a cloud-based email sending service that can be used to send transactional or marketing emails. A Lambda function can use the Amazon SES API to send an email to the user's email address after the user logs in successfully. Reference: Post authentication Lambda trigger

NEW QUESTION 6
A company is using Amazon OpenSearch Service to implement an audit monitoring system. A developer needs to create an AWS Cloudformation custom resource that is
associated with an AWS Lambda function to configure the OpenSearch Service domain. The Lambda function must access the OpenSearch Service domain by using Open Search Service internal master user credentials.
What is the MOST secure way to pass these credentials to the Lambdas function?

A. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and the Lambda function's environment variabl
B. Set the No Echo attenuate to true.
C. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and to create a paramete
D. In AWS Systems Manager Parameter Stor
E. Set the No Echo attribute to tru
F. Create an 1AM role that has the ssm GetParameter permissio
G. Assign me role to the Lambda functio
H. Store me parameter name as the Lambda function's environment variabl
I. Resolve the parameter's value at runtime.
J. Use a CloudFormation parameter to pass the master uses credentials at deployment to the OpenSearch Service domain's MasterUserOptions and the Lambda function's environment varleWe Encrypt the parameters value by using the AWS Key Management Service (AWS KMS) encrypt command.
K. Use CloudFoimalion to create an AWS Secrets Manager Secre
L. Use a CloudFormation dynamic reference to retrieve the secret's value for the OpenSearch Service domain's MasterUserOption
M. Create an 1AM role that has the secrets manage
N. GetSecretvalue permissio
O. Assign the role to the Lambda Function Store the secrets name as the Lambda function's environment variabl
P. Resole the secret's value at runtime.

Answer: D

Explanation:

The solution that will meet the requirements is to use CloudFormation to create an AWS Secrets Manager secret. Use a CloudFormation dynamic reference to retrieve the secret's value for the OpenSearch Service domain's MasterUserOptions. Create an IAM role that has the secretsmanager:GetSecretValue permission. Assign the role to the Lambda function. Store the secret's name as the Lambda function's environment variable. Resolve the secret's value at runtime. This way, the developer can pass the credentials to the Lambda function in a secure way, as AWS Secrets Manager encrypts and manages the secrets. The developer can also use a dynamic reference to avoid exposing the secret's value in plain text in the CloudFormation template. The other options either involve passing the credentials as plain text parameters, which is not secure, or encrypting them with AWS KMS, which is less convenient than using AWS Secrets Manager.
Reference: Using dynamic references to specify template values

NEW QUESTION 7

A developer maintains an Amazon API Gateway REST API. Customers use the API through a frontend UI and Amazon Cognito authentication.
The developer has a new version of the API that contains new endpoints and backward- incompatible interface changes. The developer needs to provide beta access to other developers on the team without affecting customers.
Which solution will meet these requirements with the LEAST operational overhead?

A. Define a development stage on the API Gateway AP
B. Instruct the other developers to point the endpoints to the development stage.
C. Define a new API Gateway API that points to the new API application cod
D. Instruct the other developers to point the endpoints to the new API.
E. Implement a query parameter in the API application code that determines which code version to call.
F. Specify new API Gateway endpoints for the API endpoints that the developer wants to add.

Answer: A

Explanation:

Amazon API Gateway is a service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. The developer can define a development stage on the API Gateway API and instruct the other developers to point the endpoints to the development stage. This way, the developer can provide beta access to the new version of the API without affecting customers who use the production stage. This solution will meet the requirements with the least operational overhead.
References:
? [What Is Amazon API Gateway? - Amazon API Gateway]
? [Set up a Stage in API Gateway - Amazon API Gateway]

NEW QUESTION 8

A developer needs to deploy an application running on AWS Fargate using Amazon ECS The application has environment variables that must be passed to a container for the application to initialize.
How should the environment variables be passed to the container?

A. Define an array that includes the environment variables under the environment parameter within the service definition.
B. Define an array that includes the environment variables under the environment parameter within the task definition.
C. Define an array that includes the environment variables under the entryPoint parameter within the task definition.
D. Define an array that includes the environment variables under the entryPoint parameter within the service definition.

Answer: B

Explanation:

This solution allows the environment variables to be passed to the container when it is launched by AWS Fargate using Amazon ECS. The task definition is a text file that describes one or more containers that form an application. It contains various parameters for configuring the containers, such as CPU and memory requirements, network mode, and environment variables. The environment parameter is an array of key- value pairs that specify environment variables to pass to a container. Defining an array that includes the environment variables under the entryPoint parameter within the task definition
will not pass them to the container, but use them as command-line arguments for overriding the default entry point of a container.
Defining an array that includes the environment variables under the environment or entryPoint parameter within the service definition will not pass them to the container, but cause an error because these parameters are not valid for a service definition.
Reference: [Task Definition Parameters], [Environment Variables]

NEW QUESTION 9

A developer at a company recently created a serverless application to process and show data from business reports. The application's user interface (UI) allows users to select and start processing the files. The UI displays a message when the result is available to view. The application uses AWS Step Functions with AWS Lambda functions to process the files. The developer used Amazon API Gateway and Lambda functions to create an API to support the UI.
The company's UI team reports that the request to process a file is often returning timeout errors because of the see or complexity of the files. The UI team wants the API to provide an immediate response so that the UI can deploy a message while the files are being processed. The backend process that is invoked by the API needs to send an email message when the report processing is complete.
What should the developer do to configure the API to meet these requirements?

A. Change the API Gateway route to add an X-Amz-Invocation-Type header win a sialic value of 'Event' in the integration request Deploy the API Gateway stage to apply the changes.
B. Change the configuration of the Lambda function that implements the request to process a fil
C. Configure the maximum age of the event so that the Lambda function will ion asynchronously.
D. Change the API Gateway timeout value to match the Lambda function ominous valu
E. Deploy the API Gateway stage to apply the changes.
F. Change the API Gateway route to add an X-Amz-Target header with a static value of 'A sync' in the integration request Deploy me API Gateway stage to apply the changes.

Answer: A

Explanation:

This solution allows the API to invoke the Lambda function asynchronously, which means that the API will return an immediate response without waiting for the function to complete. The X-Amz-Invocation-Type header specifies the invocation type of the Lambda function, and setting it to 'Event' means that the function

will be invoked asynchronously. The function can then use Amazon Simple Email Service (SES) to send an email message when the report processing is complete.
Reference: [Asynchronous invocation], [Set up Lambda proxy integrations in API Gateway]

NEW QUESTION 10
A developer is designing a serverless application for a game in which users register and log in through a web browser The application makes requests on behalf of users to a set of AWS Lambda functions that run behind an Amazon API Gateway HTTP API
The developer needs to implement a solution to register and log in users on the application's sign-in page. The solution must minimize operational overhead and must minimize ongoing management of user identities.
Which solution will meet these requirements'?

A. Create Amazon Cognito user pools for external social identity providers Configure 1AM roles for the identity pools.
B. Program the sign-in page to create users' 1AM groups with the 1AM roles attached to the groups
C. Create an Amazon RDS for SQL Server DB instance to store the users and manage the permissions to the backend resources in AWS
D. Configure the sign-in page to register and store the users and their passwords in an Amazon DynamoDB table with an attached IAM policy.

Answer: A

Explanation:
 https://docs.aws.amazon.com/cognito/latest/developerguide/signing-up-users-in-your-app.html

NEW QUESTION 10
An online food company provides an Amazon API Gateway HTTP API 1o receive orders for partners. The API is integrated with an AWS Lambda function. The Lambda function stores the orders in an Amazon DynamoDB table.
The company expects to onboard additional partners Some to me panthers require additional Lambda function to receive orders. The company has created an Amazon S3 bucket. The company needs 10 store all orders and updates m the S3 bucket for future analysis
How can the developer ensure that an orders and updates are stored to Amazon S3 with the LEAST development effort?

A. Create a new Lambda function and a new API Gateway API endpoin
B. Configure the new Lambda function to write to the S3 bucke
C. Modify the original Lambda function to post updates to the new API endpoint.
D. Use Amazon Kinesis Data Streams to create a new data strea
E. Modify the Lambda function to publish orders to the oats stream Configure in data stream to write to the S3 bucket.
F. Enable DynamoDB Streams on me DynamoOB tabl
G. Create a new lambda functio

H. Associate the stream's Amazon Resource Name (ARN) with the Lambda Function
bucket as records appear in the table's stream.
Configure the Lambda function to write to the S3
I. Modify the Lambda function to punish to a new Amazo
J. Simple Lambda function receives order
K. Subscribe a new Lambda function to the topi
L. Configure the new Lambda function to write to the S3 bucket as updates come through the topic.

Answer: C

Explanation:
 This solution will ensure that all orders and updates are stored to Amazon S3 with the least development effort because it uses DynamoDB Streams to capture changes in the DynamoDB table and trigger a Lambda function to write those changes to the S3 bucket. This way, the original Lambda function and API Gateway API endpoint do not need to be modified, and no additional services are required. Option A is not optimal because it will require more development effort to create a new Lambda function and a new API Gateway API endpoint, and to modify the original Lambda function to post updates to the new API endpoint. Option B is not optimal because it will introduce additional costs and complexity to use Amazon Kinesis Data Streams to create a new data stream, and to modify the Lambda function to publish orders to the data stream. Option D is not optimal because it will require more development effort to modify the Lambda function to publish to a new Amazon SNS topic, and to create and subscribe a new Lambda function to the topic. References: Using DynamoDB Streams, Using AWS Lambda with Amazon S3

NEW QUESTION 11
A developer is troubleshooting an Amazon API Gateway API Clients are receiving HTTP 400 response errors when the clients try to access an endpoint of the API.
How can the developer determine the cause of these errors?

A. Create an Amazon Kinesis Data Firehose delivery stream to receive API call logs from API Gatewa
B. Configure Amazon CloudWatch Logs as the delivery stream's destination.
C. Turn on AWS CloudTrail Insights and create a trail Specify the Amazon Resource Name (ARN) of the trail for the stage of the API.
D. Turn on AWS X-Ray for the API stage Create an Amazon CtoudWalch Logs log group Specify the Amazon Resource Name (ARN) of the log group for the API stage.
E. Turn on execution logging and access logging in Amazon CloudWatch Logs for the API stag
F. Create a CloudWatch Logs log grou
G. Specify the Amazon Resource Name (ARN) of the log group for the API stage.

Answer: D

Explanation:
 This solution will meet the requirements by using Amazon CloudWatch Logs to capture and analyze the logs from API Gateway. Amazon CloudWatch Logs is a service that monitors, stores, and accesses log files from AWS resources. The developer can turn on execution logging and access logging in Amazon CloudWatch Logs for the API stage, which enables logging information about API execution and client access to the API. The developer can create a CloudWatch Logs log group, which is a collection of log streams that share the same retention, monitoring, and access control settings. The developer can specify the Amazon Resource Name (ARN) of the log group for the API stage, which instructs API Gateway to send the logs to the specified log group. The developer can then examine the logs to determine the cause of the HTTP 400 response errors. Option A is not optimal because it will create an Amazon Kinesis Data Firehose delivery stream to receive API call logs from API Gateway, which may introduce additional costs and complexity for delivering and processing streaming data. Option B is not optimal because it will turn on AWS CloudTrail Insights and create a trail, which is a feature that helps identify and troubleshoot unusual API activity

or operational issues, not HTTP response errors. Option C is not optimal because it will turn on AWS X-Ray for the API stage, which is a service that helps analyze and debug distributed applications, not HTTP response errors. References: [Setting Up CloudWatch Logging for a REST API], [CloudWatch Logs Concepts]

NEW QUESTION 13
A company uses Amazon API Gateway to expose a set of APIs to customers. The APIs have caching enabled in API Gateway. Customers need a way to invalidate the cache for each API when they test the API.
What should a developer do to give customers the ability to invalidate the API cache?

A. Ask the customers to use AWS credentials to call the InvalidateCache API operation.
B. Attach an InvalidateCache policy to the IAM execution role that the customers use to invoke the AP
C. Ask the customers to send a request that contains the HTTP header when they make an API call.
D. Ask the customers to use the AWS SDK API Gateway class to invoke the InvalidateCache API operation.
E. Attach an InvalidateCache policy to the IAM execution role that the customers use to invoke the AP
F. Ask the customers to add the INVALIDATE_CACHE query string parameter when they make an API call.

Answer: D

NEW QUESTION 16
A developer is creating an AWS Lambda function that searches for Items from an Amazon DynamoDQ table that contains customer contact information. The DynamoDB table items have the customers as the partition and additional properties such as customer -type, name, and job_title.
The Lambda function runs whenever a user types a new character into the customer_type text Input. The developer wants to search to return partial matches of all tne email_address property of a particular customer type. The developer does not want to recreate the DynamoDB table.
What should the developer do to meet these requirements?

A. Add a global secondary index (GSI) to the DynamoDB table with customer-type input, as the partition key and email_address as the sort ke
B. Perform a query operation on the GSI by using the begins with key condition expression with the email_address property.
C.   Add a global secondary index (GSI) to the DynamoDB table with email_address as the partition key and customer_type as the sort ke
D. Perform a query operation on the GSI by using the begine_with key condition expresses with the emai
E. Address property.
F. Add a local secondary index (LSI) to the DynemoOB table with customer_type as the partition Key and email_address as the sort Ke
G. Perform a quick operation on the LSI by using the begine_with Key condition expression with the email-address property.
H. Add a local secondary index (LSI) to the DynamoDB table with job-title as the partition key and email_address as the sort ke
I. Perform a query operation on the LSI by using the begins_with key condition expression with the email_address property.

Answer: A

Explanation:
The solution that will meet the requirements is to add a global secondary index (GSI) to the DynamoDB table with customer_type as the partition key and email_address as the sort key. Perform a query operation on the GSI by using the begins_with key condition expression with the email_address property. This way, the developer can search for partial matches of the email_address property of a particular customer type without recreating the DynamoDB table. The other options either involve using a local secondary index (LSI), which requires recreating the table, or using a different partition key, which does not allow filtering by customer_type.
Reference: Using Global Secondary Indexes in DynamoDB

NEW QUESTION 20
A developer is creating a new REST API by using Amazon API Gateway and AWS Lambda. The development team tests the API and validates responses for the known use cases before deploying the API to the production environment.

The developer wants to make the REST API available for testing by using API Gateway locally. Which AWS Serverless Application Model Command Line Interface (AWS SAM CLI) subcommand will meet these requirements?

A. Sam local invoke
B. Sam local generate-event
C. Sam local start-lambda
D. Sam local start-api

Answer: D

Explanation:
? The sam local start-api subcommand allows you to run your serverless application locally for quick development and testing1. It creates a local HTTP server that acts as a proxy for API Gateway and invokes your Lambda functions based on the AWS SAM template1. You can use the sam local start-api subcommand to test your REST API locally by sending HTTP requests to the local endpoint1.

NEW QUESTION 23
A developer wants to expand an application to run in multiple AWS Regions. The developer wants to copy Amazon Machine Images (AMIs) with the latest changes and create a new application stack in the destination Region. According to company requirements, all AMIs must be encrypted in all Regions. However, not all the AMIs that the company uses are encrypted.
How can the developer expand the application to run in the destination Region while meeting the encryption requirement?

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Amazon Machine Images (AMIs) are encrypted snapshots of EC2 instances that can be used to launch new instances. The developer can create new AMIs from the existing instances and specify encryption parameters. The developer can copy the encrypted AMIs to the destination Region and use them to create a new application stack. The developer can delete the unencrypted AMIs after the encryption process is complete. This solution will meet the encryption requirement and allow the developer to expand the application to run in the destination Region.

References:
? [Amazon Machine Images (AMI) - Amazon Elastic Compute Cloud]
? [Encrypting an Amazon EBS Snapshot - Amazon Elastic Compute Cloud]
? [Copying an AMI - Amazon Elastic Compute Cloud]

NEW QUESTION 26
A company is building a compute-intensive application that will run on a fleet of Amazon EC2 instances. The application uses attached Amazon Elastic Block Store (Amazon EBS) volumes for storing data. The Amazon EBS volumes will be created at time of initial deployment. The application will process sensitive information. All of the data must be encrypted. The solution should not impact the application's performance.
Which solution will meet these requirements?

A. Configure the fleet of EC2 instances to use encrypted EBS volumes to store data.
B. Configure the application to write all data to an encrypted Amazon S3 bucket.
C. Configure a custom encryption algorithm for the application that will encrypt and decrypt all data.
D. Configure an Amazon Machine Image (AMI) that has an encrypted root volume and store the data to ephemeral disks.

Answer: A

Explanation:
 Amazon Elastic Block Store (Amazon EBS) provides block level storage volumes for use with Amazon EC2 instances1. Amazon EBS encryption offers a straight-forward encryption solution for your EBS resources associated with your EC2 instances1. When you create an encrypted EBS volume and attach it to a supported instance type, the following types of data are encrypted: Data at rest inside the volume, all data moving between the volume and the instance, all snapshots created from the volume, and all volumes created from those snapshots1. Therefore, option A is correct.

NEW QUESTION 27
A developer is creating an AWS CloudFormation template to deploy Amazon EC2 instances across multiple AWS accounts. The developer must choose the EC2 instances from a list of approved instance types.
How can the developer incorporate the list of approved instance types in the CloudFormation template?

A. Create a separate CloudFormation template for each EC2 instance type in the list.
B. In the Resources section of the CloudFormation template, create resources for each EC2 instance type in the list.
C. In the CloudFormation template, create a separate parameter for each EC2 instance type in the list.
D. In the CloudFormation template, create a parameter with the list of EC2 instance types as AllowedValues.

Answer: D

Explanation:
 In the CloudFormation template, the developer should create a parameter with the list of approved EC2 instance types as AllowedValues. This way, users can select the instance type they want to use when launching the CloudFormation stack, but only from the approved list.

NEW QUESTION 30
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? (Choose 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: AC

Explanation:
 Increasing the number of shards of the Kinesis data stream will increase the throughput and parallelism of the data processing. Increasing the memory that is allocated to the Lambda function will also increase the CPU and network performance of the function, which will reduce the run duration and improve the processing speed. Option B is not correct because decreasing the timeout of the Lambda function will not affect the processing speed, but may cause some records to fail if they exceed the timeout limit. Option D is not correct because decreasing the number of shards of the Kinesis data stream will decrease the throughput and parallelism of the data processing, which will slow down the processing speed. Option E is not correct because increasing the timeout of the Lambda function will not affect the processing speed, but may increase the cost of running the function.
References: [Amazon Kinesis Data Streams Scaling], [AWS Lambda Performance Tuning]

NEW QUESTION 35
A developer is creating an AWS Lambda function that consumes messages from an Amazon Simple Queue Service (Amazon SQS) standard queue. The developer notices that the Lambda function processes some messages multiple times.
How should developer resolve this issue MOST cost-effectively?

A. Change the Amazon SQS standard queue to an Amazon SQS FIFO queue by using the Amazon SQS message deduplication ID.
B. Set up a dead-letter queue.
C. Set the maximum concurrency limit of the AWS Lambda function to 1
D. Change the message processing to use Amazon Kinesis Data Streams instead of Amazon SQS.

Answer: A

Explanation:
 Amazon Simple Queue Service (Amazon SQS) is a fully managed queue service that allows you to de-couple and scale for applications1. Amazon SQS offers two types of queues: Standard and FIFO (First In First Out) queues1. The FIFO queue uses
the messageDeduplicationId property to treat messages with the same value as duplicate2.
Therefore, changing the Amazon SQS standard queue to an Amazon SQS FIFO queue using the Amazon SQS message deduplication ID can help resolve the

issue of the Lambda function processing some messages multiple times. Therefore, option A is correct.

NEW QUESTION 39
A developer has an application that makes batch requests directly to Amazon DynamoDB by using the BatchGetItem low-level API operation. The responses frequently return values in the UnprocessedKeys element.
Which actions should the developer take to increase the resiliency of the application when the batch response includes values in UnprocessedKeys? (Choose two.)

A. Retry the batch operation immediately.
B. Retry the batch operation with exponential backoff and randomized delay.
C. Update the application to use an AWS software development kit (AWS SDK) to make the requests.
D. Increase the provisioned read capacity of the DynamoDB tables that the operation accesses.
E. Increase the provisioned write capacity of the DynamoDB tables that the operation accesses.

Answer: BC

Explanation:
The UnprocessedKeys element indicates that the BatchGetItem operation did not process all of the requested items in the current response. This can happen if the
                                       response size limit is exceeded or if the table's provisioned throughput is exceeded. To handle this situation, the developer should retry the batch operation with exponential backoff and randomized delay to avoid throttling errors and reduce the load on the table. The developer should also use an AWS SDK to make the requests, as the SDKs automatically retry requests that return UnprocessedKeys.
References:
? [BatchGetItem - Amazon DynamoDB]
? [Working with Queries and Scans - Amazon DynamoDB]
? [Best Practices for Handling DynamoDB Throttling Errors]

NEW QUESTION 42
A developer creates a static website for their department The developer deploys the static assets for the website to an Amazon S3 bucket and serves the assets with Amazon CloudFront The developer uses origin access control (OAC) on the CloudFront distribution to access the S3 bucket
The developer notices users can access the root URL and specific pages but cannot access directories without specifying a file name. For example,
/products/index.html works, but /products returns an error The developer needs to enable accessing directories without specifying a file name without exposing the S3 bucket publicly.
Which solution will meet these requirements'?

A. Update the CloudFront distribution's settings to index.html as the default root object is set
                                       Update the Amazon S3 bucket settings and enable static website hostin
B. Specify index html as the Index document Update the S3 bucket policy to enable acces
D. Update the CloudFront distribution's origin to use the S3 website endpoint
E. Create a CloudFront function that examines the request URL and appends index.html when directories are being accessed Add the function as a viewer request CloudFront function to the CloudFront distribution's behavior.
F. Create a custom error response on the CloudFront distribution with the HTTP error code set to the HTTP 404 Not Found response code and the response page path to /index html Set the HTTP response code to the HTTP 200 OK response code

Answer: A

Explanation:
The simplest and most efficient way to enable accessing directories without specifying a file name is to update the CloudFront distribution's settings to index.html as the default root object. This will instruct CloudFront to return the index.html object when a user requests the root URL or a directory URL for the distribution. This solution does not require enabling static website hosting on the S3 bucket, creating a CloudFront function, or creating a custom error response. References
? Specifying a default root object
? cloudfront-default-root-object-configured
? How to setup CloudFront default root object?
? Ensure a default root object is configured for AWS Cloudfront …

NEW QUESTION 45
A company has an existing application that has hardcoded database credentials A developer needs to modify the existing application The application is deployed in two AWS Regions with an active-passive failover configuration to meet company's disaster recovery strategy
The developer needs a solution to store the credentials outside the code. The solution must comply With the company's disaster recovery strategy
Which solution Will meet these requirements in the MOST secure way?

A. Store the credentials in AWS Secrets Manager in the primary Regio
B. Enable secret replication to the secondary Region Update the application to use the Amazon Resource Name (ARN) based on the Region.
C. Store credentials in AWS Systems Manager Parameter Store in the primary Regio
D. Enable parameter replication to the secondary Regio
E. Update the application to use the Amazon Resource Name (ARN) based on the Region.
F. Store credentials in a config fil
G. Upload the config file to an S3 bucket in me primary Regio
H. Enable Cross-Region Replication (CRR) to an S3 bucket in the secondary regio
I. Update the application to access the config file from the S3 bucket based on the Region.
J. Store credentials in a config fil
K. Upload the config file to an Amazon Elastic File System (Amazon EFS) file syste
L. Update the application to use the Amazon EFS file system Regional endpoints to access the config file in the primary and secondary Regions.

Answer: A

Explanation:
AWS Secrets Manager is a service that allows you to store and manage secrets, such as database credentials, API keys, and passwords, in a secure and centralized way. It also provides features such as automatic secret rotation, auditing, and monitoring1. By using AWS Secrets Manager, you can avoid hardcoding credentials in your code, which is a bad security practice and makes it difficult to update them. You can also replicate your secrets to another Region, which is

useful for disaster recovery purposes2. To access your secrets from your application, you can use the ARN of the secret, which is a unique identifier that includes the Region name. This way, your application can use the appropriate secret based on the Region where it is deployed3.
References:
? AWS Secrets Manager
? Replicating and sharing secrets
? Using your own encryption keys

NEW QUESTION 49
A financial company must store original customer records for 10 years for legal reasons. A complete record contains personally identifiable information (PII). According to local regulations, PII is available to only certain people in the company and must not be shared with third parties. The company needs to make the records available to third-party organizations for statistical analysis without sharing the PII.
A developer wants to store the original immutable record in Amazon S3. Depending on who accesses the S3 document, the document should be returned as is or with all the PII removed. The developer has written an AWS Lambda function to remove the PII from the document. The function is named removePii.
What should the developer do so that the company can meet the PII requirements while maintaining only one copy of the document?

A. Set up an S3 event notification that invokes the removePii function when an S3 GET request is mad
B. Call Amazon S3 by using a GET request to access the object without PII.
C. Set up an S3 event notification that invokes the removePii function when an S3 PUT request is mad
D. Call Amazon S3 by using a PUT request to access the object without PII.
E. Create an S3 Object Lambda access point from the S3 consol
F. Select the removePii functio
G. Use S3 Access Points to access the object without PII.
H. Create an S3 access point from the S3 consol
I. Use the access point name to call the GetObjectLegalHold S3 API functio
J. Pass in the removePii function name to access the object without PII.

Answer: C

Explanation:
 S3 Object Lambda allows you to add your own code to process data retrieved from S3 before returning it to an application. You can use an AWS Lambda function to modify the data, such as removing PII, redacting confidential information, or resizing images. You can create an S3 Object Lambda access point and associate it with your Lambda function. Then, you can use the access point to request objects from S3 and get the modified data back. This way, you can maintain only one copy of the original

document in S3 and apply different transformations depending on who accesses it. Reference: Using AWS Lambda with Amazon S3

NEW QUESTION 54
For a deployment using AWS Code Deploy, what is the run order of the hooks for in-place deployments?

A. BeforeInstall -> ApplicationStop -> ApplicationStart -> AfterInstall
B. ApplicationStop -> BeforeInstall -> AfterInstall -> ApplicationStart
C. BeforeInstall -> ApplicationStop -> ValidateService -> ApplicationStart
D. ApplicationStop -> BeforeInstall -> ValidateService -> ApplicationStart

Answer: B

Explanation:
For in-place deployments, AWS CodeDeploy uses a set of predefined hooks that run in a specific order during each deployment lifecycle event. The hooks are ApplicationStop, BeforeInstall, AfterInstall, ApplicationStart, and ValidateService. The run order of the hooks for in-place deployments is as follows:
? ApplicationStop: This hook runs first on all instances and stops the current
application that is running on the instances.
? BeforeInstall: This hook runs after ApplicationStop on all instances and performs any tasks required before installing the new application revision.
? AfterInstall: This hook runs after BeforeInstall on all instances and performs any
tasks required after installing the new application revision.
? ApplicationStart: This hook runs after AfterInstall on all instances and starts the new application that has been installed on the instances.
? ValidateService: This hook runs last on all instances and verifies that the new application is running properly on the instances.
Reference: [AWS CodeDeploy lifecycle event hooks reference]

NEW QUESTION 56
A developer is building a serverless application by using AWS Serverless Application Model (AWS SAM) on multiple AWS Lambda functions. When the application is deployed, the developer wants to shift 10% of the traffic to the new deployment of the application for the first 10 minutes after deployment. If there are no issues, all traffic must switch over to the new version.
Which change to the AWS SAM template will meet these requirements?

A. Set the Deployment Preference Type to Canaryl OPercent10Minute
B. Set the AutoPublishAlias property to the Lambda alias.
C. Set the Deployment Preference Type to Linearl OPercentEveryIOMinute
D. Set AutoPublishAlias property to the Lambda alias.
E. Set the Deployment Preference Type to Canaryl OPercentIOMinute
F. Set the PreTraffic and PostTraffic properties to the Lambda alias.
G. Set the Deployment Preference Type to Linearl OPercentEvery10Minute
H. Set PreTraffic and PostTraffic properties to the Lambda alias.

Answer: A

Explanation:
? The Deployment Preference Type property specifies how traffic should be shifted between versions of a Lambda function1. The Canary10Percent10Minutes option means that 10% of the traffic is immediately shifted to the new version, and after 10 minutes, the remaining 90% of the traffic is shifted1. This matches the requirement of shifting 10% of the traffic for the first 10 minutes, and then switching all traffic to the new version.
? The AutoPublishAlias property enables AWS SAM to automatically create and update a Lambda alias that points to the latest version of the function1. This is

required to use the Deployment Preference Type property1. The alias name can be specified by the developer, and it can be used to invoke the function with the latest code.

NEW QUESTION 58
A developer has an application that is composed of many different AWS Lambda functions. The Lambda functions all use some of the same dependencies. To avoid security issues the developer is constantly updating the dependencies of all of the Lambda functions. The result is duplicated effort to reach function.
How can the developer keep the dependencies of the Lambda functions up to date with the LEAST additional complexity?

A. Define a maintenance window for the Lambda functions to ensure that the functions get updated copies of the dependencies.
B. Upgrade the Lambda functions to the most recent runtime version.
C. Define a Lambda layer that contains all of the shared dependencies.
D. Use an AWS CodeCommit repository to host the dependencies in a centralized location.

Answer: C

Explanation:
This solution allows the developer to keep the dependencies of the Lambda functions up to date with the least additional complexity because it eliminates the need to update each function individually. A Lambda layer is a ZIP archive that contains libraries, custom runtimes, or other dependencies. The developer can create a layer that contains all of the shared dependencies and attach it to multiple Lambda functions. When the developer updates the layer, all of the functions that use the layer will have access to the latest version of the dependencies.
Reference: [AWS Lambda layers]

NEW QUESTION 61
A company's developer has deployed an application in AWS by using AWS CloudFormation The CloudFormation stack includes parameters in AWS Systems Manager Parameter Store that the application uses as configuration settings. The application can modify the parameter values
When the developer updated the stack to create additional resources with tags, the developer noted that the parameter values were reset and that the values ignored the latest changes made by the application. The developer needs to change the way the company deploys the CloudFormation stack. The developer also needs to avoid resetting the parameter values outside the stack.
Which solution will meet these requirements with the LEAST development effort?

A. Modify the CloudFormation stack to set the deletion policy to Retain for the Parameter Store parameters.
B. Create an Amazon DynamoDB table as a resource in the CloudFormation stack to hold configuration data for the application Migrate the parameters that the application is modifying from Parameter Store to the DynamoDB table
C. Create an Amazon RDS DB instance as a resource in the CloudFormation stac
D. Create a table in the database for parameter configuratio
E. Migrate the parameters that the application is modifying from Parameter Store to the configuration table
F. Modify the CloudFormation stack policy to deny updates on Parameter Store parameters

Answer: D

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/protect-stack- resources.html#stack-policy-samples

NEW QUESTION 66
A developer uses AWS CloudFormation to deploy an Amazon API Gateway API and an AWS Step Functions state machine The state machine must reference the API Gateway API after the CloudFormation template is deployed The developer needs a solution that uses the state machine to reference the API Gateway endpoint.
Which solution will meet these requirements MOST cost-effectively?

A. Configure the CloudFormation template to reference the API endpoint in the DefinitionSubstitutions property for the AWS StepFunctions StateMachme resource.
B. Configure the CloudFormation template to store the API endpoint in an environment variable for the AWS::StepFunctions::StateMachine resourc Configure the state machine to reference the environment variable
C. Configure the CloudFormation template to store the API endpoint in a standard AWS: SecretsManager Secret resource Configure the state machine to reference the resource
D. Configure the CloudFormation template to store the API endpoint in a standard AWS::AppConfig;:ConfigurationProfile resource Configure the state machine to reference the resource.

Answer: A

Explanation:
The most cost-effective solution is to use the DefinitionSubstitutions property of the AWS::StepFunctions::StateMachine resource to inject the API endpoint as a variable in the state machine definition. This way, the developer can use the intrinsic function
Fn::GetAtt to get the API endpoint from the AWS::ApiGateway::RestApi resource, and pass it to the state machine without creating any additional resources or environment variables. The other solutions involve creating and managing extra resources, such as Secrets Manager secrets or AppConfig configuration profiles, which incur additional costs and complexity. References
? AWS::StepFunctions::StateMachine - AWS CloudFormation
? Call API Gateway with Step Functions - AWS Step Functions
? amazon-web-services aws-api-gateway terraform aws-step-functions

NEW QUESTION 69
A company runs a batch processing application by using AWS Lambda functions and Amazon API Gateway APIs with deployment stages for development, user acceptance testing and production A development team needs to configure the APIs in the deployment stages to connect to third-party service endpoints.
Which solution will meet this requirement?

A. Store the third-party service endpoints in Lambda layers that correspond to the stage
B. Store the third-party service endpoints in API Gateway stage variables that correspond to the stage

C. Encode the third-party service endpoints as query parameters in the API Gateway request URL.
D. Store the third-party service endpoint for each environment in AWS AppConfig

Answer: B

Explanation:
API Gateway stage variables are name-value pairs that can be defined as configuration attributes associated with a deployment stage of a REST API. They act like environment variables and can be used in the API setup and mapping templates. For example, the development team can define a stage variable named endpoint and assign it different values for each stage, such as dev.example.com for development, uat.example.com for user acceptance testing, and prod.example.com for production. Then, the team can use the stage variable value in the integration request URL, such as http://$ { stageVariables.endpoint}/api. This way, the team can use the same API setup with different endpoints at each stage by resetting the stage variable value. The other solutions are either not feasible or not cost-effective. Lambda layers are used to package and load dependencies for Lambda functions, not for storing endpoints. Encoding the endpoints as query parameters would expose them to the public and make the request URL unnecessarily long. Storing the endpoints in AWS AppConfig would incur additional costs and complexity, and would require additional logic to retrieve the values from the configuration store. References
? Using Amazon API Gateway stage variables
? Setting up stage variables for a REST API deployment
? Setting stage variables using the Amazon API Gateway console

NEW QUESTION 72
A company is using Amazon RDS as the Backend database for its application. After a recent marketing campaign, a surge of read requests to the database increased the latency of data retrieval from the database.
The company has decided to implement a caching layer in front of the database. The cached content must be encrypted and must be highly available.
Which solution will meet these requirements?

A. Amazon Cloudfront
B. Amazon ElastiCache to Memcached
C. Amazon ElastiCache for Redis in cluster mode
D. Amazon DynamoDB Accelerate (DAX)

Answer: C

Explanation:
This solution meets the requirements because it provides a caching layer that can store and retrieve encrypted data from multiple nodes. Amazon ElastiCache for Redis supports encryption at rest and in transit, and can scale horizontally to increase the cache capacity and availability. Amazon ElastiCache for Memcached does not support encryption, Amazon CloudFront is a content delivery network that is not suitable for caching database queries, and Amazon DynamoDB Accelerator (DAX) is a caching service that only works with DynamoDB tables.
Reference: [Amazon ElastiCache for Redis Features], [Choosing a Cluster Engine]

NEW QUESTION 76
A developer is working on an ecommerce website The developer wants to review server logs without logging in to each of the application servers individually. The website runs on multiple Amazon EC2 instances, is written in Python, and needs to be highly available
How can the developer update the application to meet these requirements with MINIMUM changes?

A. Rewrite the application to be cloud native and to run on AWS Lambda, where the logs can be reviewed in Amazon CloudWatch
B. Set up centralized logging by using Amazon OpenSearch Service, Logstash, and OpenSearch Dashboards
C. Scale down the application to one larger EC2 instance where only one instance is recording logs
D. Install the unified Amazon CloudWatch agent on the EC2 instances Configure the agent to push the application logs to CloudWatch

Answer: D

Explanation:
The unified Amazon CloudWatch agent can collect both system metrics and log files from Amazon EC2 instances and on-premises servers. By installing and configuring the agent on the EC2 instances, the developer can easily access and analyze the application logs in CloudWatch without logging in to each server individually. This option requires minimum changes to the existing application and does not affect its availability or scalability. References
? Using the CloudWatch Agent
? Collecting Metrics and Logs from Amazon EC2 Instances and On-Premises Servers with the CloudWatch Agent

NEW QUESTION 77
A developer is building a web application that uses Amazon API Gateway to expose an AWS Lambda function to process requests from clients. During testing, the developer notices that the API Gateway times out even though the Lambda function finishes under the set time limit.
Which of the following API Gateway metrics in Amazon CloudWatch can help the developer troubleshoot the issue? (Choose two.)

A. CacheHitCount
B. IntegrationLatency
C. CacheMissCount
D. Latency
E. Count

Answer: BD

Explanation:
Amazon API Gateway is a service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. Amazon CloudWatch is a service that monitors AWS resources and applications. API Gateway provides several CloudWatch metrics to help developers troubleshoot issues with their APIs. Two of the metrics that can help the developer troubleshoot the issue of API Gateway timing out are:
? IntegrationLatency: This metric measures the time between when API Gateway relays a request to the backend and when it receives a response from the backend. A high value for this metric indicates that the backend is taking too long to respond and may cause API Gateway to time out.
? Latency: This metric measures the time between when API Gateway receives a request from a client and when it returns a response to the client. A high value for this metric indicates that either the integration latency is high or API Gateway is

taking too long to process the request or response.
References:
? [What Is Amazon API Gateway? - Amazon API Gateway]
? [Amazon API Gateway Metrics and Dimensions - Amazon CloudWatch]
? [Troubleshooting API Errors - Amazon API Gateway]

NEW QUESTION 82
A developer is designing a serverless application with two AWS Lambda functions to process photos. One Lambda function stores objects in an Amazon S3 bucket and stores the associated metadata in an Amazon DynamoDB table. The other Lambda function fetches the objects from the S3 bucket by using the metadata from the DynamoDB table. Both Lambda functions use the same Python library to perform complex computations and are approaching the quota for the maximum size of zipped deployment packages.
What should the developer do to reduce the size of the Lambda deployment packages with the LEAST operational overhead?

A. Package each Python library in its own .zip file archiv
B. Deploy each Lambda function with its own copy of the library.
C. Create a Lambda layer with the required Python librar
D. Use the Lambda layer in both Lambda functions.
E. Combine the two Lambda functions into one Lambda functio
F. Deploy the Lambda function as a single .zip file archive.
G. Download the Python library to an S3 bucke
H. Program the Lambda functions to reference the object URLs.

Answer: B

Explanation:
 AWS Lambda is a service that lets developers run code without provisioning or managing servers. Lambda layers are a distribution mechanism for libraries, custom runtimes, and other dependencies. The developer can create a Lambda layer with the

required Python library and use the layer in both Lambda functions. This will reduce the size of the Lambda deployment packages and avoid reaching the quota for the maximum size of zipped deployment packages. The developer can also benefit from using layers to manage dependencies separately from function code.
References:
? [What Is AWS Lambda? - AWS Lambda]
? [AWS Lambda Layers - AWS Lambda]

NEW QUESTION 87
A company developed an API application on AWS by using Amazon CloudFront, Amazon API Gateway, and AWS Lambda. The API has a minimum of four requests every second. A developer notices that many API users run the same query by using the POST method. The developer wants to cache the POST request to optimize the API resources. Which solution will meet these requirements?

A.

Configure the CloudFront cach

B. Update the application to return cached content based upon the default request headers.

C. Override the cache method in the selected stage of API Gatewa

D. Select the POST method.

E. Save the latest request response in Lambda /tmp director

F. Update the Lambda function to check the /tmp directory.

G. Save the latest request in AWS Systems Manager Parameter Stor

H. Modify the Lambda function to take the latest request response from Parameter Store.

Answer: B

Explanation:

Amazon API Gateway provides tools for creating and documenting web APIs that route HTTP requests to Lambda functions2. You can secure access to your API with authentication and authorization controls. Your APIs can serve traffic over the internet or can be accessible only within your VPC2. You can override the cache method in the selected stage of API Gateway2. Therefore, option B is correct.

NEW QUESTION 92

A company is planning to use AWS CodeDeploy to deploy an application to Amazon Elastic Container Service (Amazon ECS) During the deployment of a new version of the application, the company initially must expose only 10% of live traffic to the new version of the deployed application. Then, after 15 minutes elapse, the company must route all the remaining live traffic to the new version of the deployed application.

Which CodeDeploy predefined configuration will meet these requirements?

A. CodeDeployDefault ECSCanary10Percent15Minutes

B. CodeDeployDefault LambdaCanary10Percent5Minutes

C. CodeDeployDefault LambdaCanary10Percent15Minutes

D. CodeDeployDefault ECSLinear10PercentEvery1 Minutes

Answer: A

Explanation:

The predefined configuration "CodeDeployDefault.ECSCanary10Percent15Minutes" is designed for Amazon Elastic Container Service (Amazon ECS) deployments and meets the specified requirements. It will perform a canary deployment, which means it will initially route 10% of live traffic to the new version of the application, and then after 15 minutes elapse, it will automatically route all the remaining live traffic to the new version. This gradual deployment approach allows

the company to verify the health and performance of the new version with a small portion of traffic before fully deploying it to all

users.

NEW QUESTION 96
An ecommerce application is running behind an Application Load Balancer. A developer observes some unexpected load on the application during non-peak hours. The developer wants to analyze patterns for the client IP addresses that use the application. Which HTTP header should the developer use for this analysis?

A. The X-Forwarded-Proto header
B. The X-F Forwarded-Host header
C. The X-Forwarded-For header
D. The X-Forwarded-Port header

Answer: C

Explanation:
The HTTP header that the developer should use for this analysis is the X- Forwarded-For header. This header contains the IP address of the client that made the request to the Application Load Balancer. The developer can use this header to analyze patterns for the client IP addresses that use the application. The other headers either contain information about the protocol, host, or port of the request, which are not relevant for the analysis.
Reference: How Application Load Balancer works with your applications

NEW QUESTION 100
A company developed an API application on AWS by using Amazon CloudFront. Amazon API Gateway, and AWS Lambda. The API has a minimum of four requests every second A developer notices that many API users run the same query by using the POST method. The developer wants to cache the POST request to optimize the API resources.
Which solution will meet these requirements'?

A. Configure the CloudFront cache Update the application to return cached content based upon the default request headers.
B. Override the cache method in me selected stage of API Gateway Select the POST method.
C. Save the latest request response in Lambda /tmp directory Update the Lambda function to check the /tmp directory
D. Save the latest request m AWS Systems Manager Parameter Store Modify the Lambda function to take the latest request response from Parameter Store

Answer: A

Explanation:
This solution will meet the requirements by using Amazon CloudFront, which is a content delivery network (CDN) service that speeds up the delivery of web content and APIs to end users. The developer can configure the CloudFront cache, which is a set of edge locations that store copies of popular or recently accessed content close to the viewers. The developer can also update the application to return cached content based upon the default request headers, which are a set of HTTP headers that CloudFront automatically forwards to the origin server and uses to determine whether an object in an edge location is still valid. By caching the POST requests, the developer can optimize the API resources and reduce the latency for repeated queries. Option B is not optimal because it will override the cache method in the selected stage of API Gateway, which is not possible or effective as API Gateway does not support caching for POST methods by default. Option C is not optimal because it will save the latest request response in Lambda /tmp directory, which is a local storage space that is available for each Lambda function invocation, not a cache that can be shared across multiple invocations or requests. Option D is not optimal because it will save the latest request in AWS Systems Manager Parameter Store, which is a service that provides secure and scalable storage for configuration data and secrets, not a cache for API responses.
References: [Amazon CloudFront], [Caching Content Based on Request Headers]

NEW QUESTION 101
An application uses an Amazon EC2 Auto Scaling group. A developer notices that EC2 instances are taking a long time to become available during scale-out events. The UserData script is taking a long time to run.
The developer must implement a solution to decrease the time that elapses before an EC2 instance becomes available. The solution must make the most recent version of the application available at all times and must apply all available security updates. The solution also must minimize the number of images that are created. The images must be validated.
Which combination of steps should the developer take to meet these requirements? (Choose two.)

A. Use EC2 Image Builder to create an Amazon Machine Image (AMI). Install all the patches and agents that are needed to manage and run the applicatio

B. Update the Auto Scaling group launch configuration to use the AMI.

C. Use EC2 Image Builder to create an Amazon Machine Image (AMI). Install the latest version of the application and all the patches and agents that are needed to manage and run the applicatio

D. Update the Auto Scaling group launch configuration to use the AMI.

E. Set up AWS CodeDeploy to deploy the most recent version of the application at runtime.

F. Set up AWS CodePipeline to deploy the most recent version of the application at runtime.

G. Remove any commands that perform operating system patching from the UserData script.

Answer: BE

Explanation:

 AWS CloudFormation is a service that enables developers to model and provision AWS resources using templates. The developer can use the following steps to avoid accidental database deletion in the future:

? Set up AWS CodeDeploy to deploy the most recent version of the application at

runtime. This will ensure that the application code is always up to date and does not depend on the AMI.

? Remove any commands that perform operating system patching from the

UserData script. This will reduce the time that the UserData script takes to run and speed up the instance launch process.

References:

? [What Is AWS CloudFormation? - AWS CloudFormation]

? [What Is AWS CodeDeploy? - AWS CodeDeploy]

? [Running Commands on Your Linux Instance at Launch - Amazon Elastic Compute Cloud]

NEW QUESTION 106

A developer is writing a serverless application that requires an AWS Lambda function to be invoked every 10 minutes.

What is an automated and serverless way to invoke the function?

A. Deploy an Amazon EC2 instance based on Linux, and edit its /etc/confab file by adding a command to periodically invoke the lambda function

B. Configure an environment variable named PERIOD for the Lambda functio

C. Set the value to 600.

D. Create an Amazon EventBridge rule that runs on a regular schedule to invoke the Lambda function.

E. Create an Amazon Simple Notification Service (Amazon SNS) topic that has a subscription to the Lambda function with a 600-second timer.

Answer: C

Explanation:

 The solution that will meet the requirements is to create an Amazon EventBridge rule that runs on a regular schedule to invoke the Lambda function. This way, the developer can use an automated and serverless way to invoke the function every 10 minutes. The developer can also use a cron expression or a rate expression to specify the schedule for the rule. The other options either involve using an Amazon EC2 instance, which is not serverless, or using environment variables or query parameters, which do not trigger the function.

Reference: Schedule AWS Lambda functions using EventBridge

NEW QUESTION 108
A company has an Amazon S3 bucket containing premier content that it intends to make available to only paid subscribers of its website. The S3 bucket currently has default permissions of all objects being private to prevent inadvertent exposure of the premier content to non-paying website visitors.
How can the company Limit the ability to download a premier content file in the S3 Bucket to paid subscribers only?

A. Apply a bucket policy that allows anonymous users to download the content from the S3 bucket.
B. Generate a pre-signed object URL for the premier content file when a pad subscriber requests a download.
C. Add a Docket policy that requires multi-factor authentication for request to access the S3 bucket objects.
D. Enable server-side encryption on the S3 bucket for data protection against the non- paying website visitors.

Answer: B

Explanation:
This solution will limit the ability to download a premier content file in the S3 bucket to paid subscribers only because it uses a pre-signed object URL that grants temporary access to an S3 object for a specified duration. The pre-signed object URL can be generated by the company's website when a paid subscriber requests a download, and can be verified by Amazon S3 using the signature in the URL. Option A is not optimal because it will allow anyone to download the content from the S3 bucket without verifying their subscription status. Option C is not optimal because it will require additional steps and costs to configure multi-factor authentication for accessing the S3 bucket objects, which may not be feasible or user-friendly for paid subscribers. Option D is not optimal because it will not prevent non-paying website visitors from accessing the S3 bucket objects, but only encrypt them at rest.
References: Share an Object with Others, [Using Amazon S3 Pre-Signed URLs]

NEW QUESTION 111
When using the AWS Encryption SDK how does the developer keep track of the data encryption keys used to encrypt data?

A. The developer must manually keep Hack of the data encryption keys used for each data object.
B. The SDK encrypts the data encryption key and stores it (encrypted) as part of the resumed ophertext.
C. The SDK stores the data encryption keys automaticity in Amazon S3.
D. The data encryption key is stored m the user data for the EC2 instance.

Answer: B

Explanation:
This solution will meet the requirements by using AWS Encryption SDK, which is a client-side encryption library that enables developers to encrypt and decrypt data using data encryption keys that are protected by AWS Key Management Service (AWS KMS). The SDK encrypts the data encryption key with a customer master key (CMK) that is managed by AWS KMS, and stores it (encrypted) as part of the returned ciphertext. The developer does not need to keep track of the data encryption keys used to encrypt data, as they are stored with the encrypted data and can be retrieved and decrypted by using AWS KMS when needed. Option A is not optimal because it will require manual tracking of the data encryption keys used for each data object, which is error-prone and inefficient. Option C is not optimal because it will store the data encryption keys automatically in Amazon S3, which is unnecessary and insecure as Amazon S3 is not designed for storing encryption keys. Option D is not optimal because it will store the data encryption key in the user data for the EC2 instance, which is also unnecessary and insecure as user data is not encrypted by default.
References: [AWS Encryption SDK], [AWS Key Management Service]

NEW QUESTION 116
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 m 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: C

Explanation:
This step should be completed prior to deploying the application because it prepares the application artifacts for deployment. The AWS Serverless Application Model (AWS SAM) is a framework that simplifies building and deploying serverless applications on AWS. The AWS SAM CLI is a command-line tool that helps you create, test, and deploy serverless applications using AWS SAM templates. The sam package command bundles the application artifacts, such as Lambda function code and API definitions, and uploads them to an Amazon S3 bucket. The command also returns a CloudFormation template that is ready to be deployed with the sam deploy command. Compressing the application to a zip file and uploading it to AWS Lambda will not work because it does not use AWS SAM templates or CloudFormation. Testing the new Lambda function by first tracing it in AWS X- Ray will not prepare the application for deployment, but only monitor its performance and errors. Creating the application environment using the eb create my-env command will not work because it is a command for AWS Elastic Beanstalk, not AWS SAM.

NEW QUESTION 119
An application uses Lambda functions to extract metadata from files uploaded to an S3 bucket; the metadata is stored in Amazon DynamoDB. The application starts behaving unexpectedly, and the developer wants to examine the logs of the Lambda function code for errors.
Based on this system configuration, where would the developer find the logs?

A. Amazon S3
B. AWS CloudTrail
C. Amazon CloudWatch
D. Amazon DynamoDB

Answer: C

Explanation:
Amazon CloudWatch is the service that collects and stores logs from AWS Lambda functions. The developer can use CloudWatch Logs Insights to query and analyze the logs for errors and metrics. Option A is not correct because Amazon S3 is a storage service that does not store Lambda function logs. Option B is not correct because AWS CloudTrail is a service that records API calls and events for AWS services, not Lambda function logs. Option D is not correct because Amazon DynamoDB is a database service that does not store Lambda function logs.
References: AWS Lambda Monitoring, [CloudWatch Logs Insights]

NEW QUESTION 122
A developer is building a new application on AWS. The application uses an AWS Lambda function that retrieves information from an Amazon DynamoDB table. The developer hard coded the DynamoDB table name into the Lambda function code. The table name might change over time. The developer does not want to modify the Lambda code if the table name changes.
Which solution will meet these requirements MOST efficiently?

A. Create a Lambda environment variable to store the table nam
B. Use the standard method for the programming language to retrieve the variable.
C. Store the table name in a fil
D. Store the file in the /tmp folde
E. Use the SDK for the programming language to retrieve the table name.
F. Create a file to store the table nam

G. Zip the file and upload the file to the Lambda laye
H. Use the SDK for the programming language to retrieve the table name.
          Create a global variable that is outside the handler in the Lambda function to store the table name.
I.

Answer: A

Explanation:
The solution that will meet the requirements most efficiently is to create a Lambda environment variable to store the table name. Use the standard method for the programming language to retrieve the variable. This way, the developer can avoid hard- coding the table name in the Lambda function code and easily change the table name by updating the environment variable. The other options either involve storing the table name in a file, which is less efficient and secure than using an environment variable, or creating a global variable, which is not recommended as it can cause concurrency issues.
Reference: Using AWS Lambda environment variables

NEW QUESTION 127
A developer is using an AWS Lambda function to generate avatars for profile pictures that are uploaded to an Amazon S3 bucket. The Lambda function is automatically invoked for profile pictures that are saved under the /original/ S3 prefix. The developer notices that some pictures cause the Lambda function to time out. The developer wants to implement a fallback mechanism by using another Lambda function that resizes the profile picture.
          Which solution will meet these requirements with the LEAST development effort?

A. Set the image resize Lambda function as a destination of the avatar generator Lambda function for the events that fail processing.
B. Create an Amazon Simple Queue Service (Amazon SQS) queu
C. Set the SQS queue as a destination with an on failure condition for the avatar generator Lambda functio
D. Configure the image resize Lambda function to poll from the SQS queue.
E. Create an AWS Step Functions state machine that invokes the avatar generator Lambda function and uses the image resize Lambda function as a fallbac
F. Create an Amazon EventBridge rule that matches events from the S3 bucket to invoke the state machine.
G. Create an Amazon Simple Notification Service (Amazon SNS) topi
H. Set the SNS topic as a destination with an on failure condition for the avatar generator Lambda functio
I. Subscribe the image resize Lambda function to the SNS topic.

Answer: A

Explanation:
The solution that will meet the requirements with the least development effort is to set the image resize Lambda function as a destination of the avatar generator Lambda function for the events that fail processing. This way, the fallback mechanism is automatically triggered by the Lambda service without requiring any additional components or configuration. The other options involve creating and managing additional resources such as queues, topics, state machines, or rules, which would increase the complexity and cost of the solution.
Reference: Using AWS Lambda destinations

NEW QUESTION 131
A company hosts a client-side web application for one of its subsidiaries on Amazon S3. The web application can be accessed through Amazon CloudFront from https://www.example.com. After a successful rollout, the company wants to host three more client-side web applications for its remaining subsidiaries on three separate S3 buckets.
To achieve this goal, a developer moves all the common JavaScript files and web fonts to a central S3 bucket that serves the web applications. However, during testing, the developer notices that the browser blocks the JavaScript files and web fonts.
What should the developer do to prevent the browser from blocking the JavaScript files and web fonts?

A. Create four access points that allow access to the central S3 bucke
B. Assign an access point to each web application bucket.
C. Create a bucket policy that allows access to the central S3 bucke
D. Attach the bucket policy to the central S3 bucket.
E. Create a cross-origin resource sharing (CORS) configuration that allows access to the central S3 bucke
F. Add the CORS configuration to the central S3 bucket.
G. Create a Content-MD5 header that provides a message integrity check for the central S3 bucke
H. Insert the Content-MD5 header for each web application request.

Answer: C

Explanation:
This is a frequent trouble. Web applications cannot access the resources in other domains by default, except some exceptions. You must configure CORS on the resources to be accessed. https://docs.aws.amazon.com/AmazonS3/latest/userguide/cors.html

NEW QUESTION 135
A company has multiple Amazon VPC endpoints in the same VPC. A developer needs configure an Amazon S3 bucket policy so users can access an S3 bucket only by using these VPC endpoints.
Which solution will meet these requirements?

A. Create multiple S3 bucket polices by using each VPC endpoint ID that have the aws SourceVpce value in the StringNotEquals condition.
B. Create a single S3 bucket policy that has the aws SourceVpc value and in the StingNotEquals condition to use VPC ID.
C. Create a single S3 bucket policy that the multiple aws SourceVpce value and in the SringNotEquals condton to use vpce.
D. Create a single S3 bucket policy that has multiple aws sourceVpce value in the StingNotEquale conditio
E. Repeat for all the VPC endpoint IDs.

Answer: D

Explanation:
This solution will meet the requirements by creating a single S3 bucket policy that denies access to the S3 bucket unless the request comes from one of the specified VPC endpoints. The aws:SourceVpce condition key is used to match the ID of the VPC endpoint that is used to access the S3 bucket. The
                                                                                allowed.
StringNotEquals condition operator is used to negate the condition, so that only requests from the listed VPC endpoints are

Option A is not optimal because it will create multiple S3 bucket policies, which is not possible as only one bucket policy can be attached to an S3 bucket. Option B is not optimal because it will use the aws:SourceVpc condition key, which matches the ID of the VPC that is used to access the S3 bucket, not the VPC endpoint. Option C is not optimal because it will use the StringNotEquals condition operator with a single value, which will deny access to the S3 bucket from all VPC endpoints except one.
References: Using Amazon S3 Bucket Policies and User Policies, AWS Global Condition Context Keys

NEW QUESTION 136
A company is implementing an application on Amazon EC2 instances. The application needs to process incoming transactions. When the application detects a transaction that is not valid, the application must send a chat message to the company's support team. To send the message, the application needs to retrieve the access token to authenticate by using the chat API.
A developer needs to implement a solution to store the access token. The access token must be encrypted at rest and in transit. The access token must also be accessible from other AWS accounts.
Which solution will meet these requirements with the LEAST management overhead?

A. Use an AWS Systems Manager Parameter Store SecureString parameter that uses an AWS Key Management Service (AWS KMS) AWS managed key to store the access toke
B. Add a resource-based policy to the parameter to allow access from other account
C. Update the IAM role of the EC2 instances with permissions to access Parameter Stor
D. Retrieve the token from Parameter Store with the decrypt flag enable
E. Use the decrypted access token to send the message to the chat.
F. Encrypt the access token by using an AWS Key Management Service (AWS KMS) customer managed ke
G. Store the access token in an Amazon DynamoDB tabl
H. Update the IAM role of the EC2 instances with permissions to access DynamoDB and AWS KM
I. Retrieve the token from DynamoD
J. Decrypt the token by using AWS KMS on the EC2 instance
K. Use the decrypted access token to send the message to the chat.
L. Use AWS Secrets Manager with an AWS Key Management Service (AWS KMS) customer managed key to store the access toke
M. Add a resource-based policy to the secret to allow access from other account
N. Update the IAM role of the EC2 instances with permissions to access Secrets Manage
O. Retrieve the token from Secrets Manage
P. Use the decrypted access token to send the message to the chat.
Q. Encrypt the access token by using an AWS Key Management Service (AWS KMS) AWS managed ke
R. Store the access token in an Amazon S3 bucke
S. Add a bucket policy to the S3 bucket to allow access from other account
T. Update the IAM role of the EC2 instances with permissions to access Amazon S3 and AWS KM
. Retrieve the token from the S3 bucke
. Decrypt the token by using AWS KMS on the EC2 instance
. Use the decrypted access token to send the massage to the chat.

Answer: C

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/secrets-manager-share-between-accounts/
https://docs.aws.amazon.com/secretsmanager/latest/userguide/auth-and- access_examples_cross.html

NEW QUESTION 140
A company runs an application on AWS The application stores data in an Amazon DynamoDB table Some queries are taking a long time to run These slow queries involve an attribute that is not the table's partition key or sort key
The amount of data that the application stores in the DynamoDB table is expected to increase significantly. A developer must increase the performance of the queries.
Which solution will meet these requirements'?

A. Increase the page size for each request by setting the Limit parameter to be higher than the default value Configure the application to retry any request that exceeds the provisioned throughput.
B. Create a global secondary index (GSI). Set query attribute to be the partition key of the index
C. Perform a parallel scan operation by issuing individual scan requests in the parameters specify the segment for the scan requests and the total number of segments for the parallel scan.
D. Turn on read capacity auto scaling for the DynamoDB tabl
E. Increase the maximum read capacity units (RCUs).

Answer: B

Explanation:
Creating a global secondary index (GSI) is the best solution to improve the performance of the queries that involve an attribute that is not the table's partition key or sort key. A GSI allows you to define an alternate key for your table and query the data using that key. This way, you can avoid scanning the entire table and reduce the latency and cost of your queries. You should also follow the best practices for designing and using GSIs in DynamoDB12. References
? Working with Global Secondary Indexes - Amazon DynamoDB
? DynamoDB Performance & Latency - Everything You Need To Know

NEW QUESTION 142
A developer is migrating an application to Amazon Elastic Kubernetes Service (Amazon EKS). The developer migrates the application to Amazon Elastic Container Registry (Amazon ECR) with an EKS cluster.
As part of the application migration to a new backend, the developer creates a new AWS account. The developer makes configuration changes to the application to point the application to the new AWS account and to use new backend resources. The developer successfully tests the changes within the application by deploying the pipeline.
The Docker image build and the pipeline deployment are successful, but the application is still connecting to the old backend. The developer finds that the application's configuration is still referencing the original EKS cluster and not referencing the new backend resources.
Which reason can explain why the application is not connecting to the new resources?

A. The developer did not successfully create the new AWS account.

B. The developer added a new tag to the Docker image.
C. The developer did not update the Docker image tag to a new version.
D. The developer pushed the changes to a new Docker image tag.

Answer: C

Explanation:
The correct answer is C. The developer did not update the Docker image tag to a new version.
* C. The developer did not update the Docker image tag to a new version. This is correct. When deploying an application to Amazon EKS, the developer needs to specify the Docker image tag that contains the application code and configuration. If the developer does not update the Docker image tag to a new version after making changes to the application, the EKS cluster will continue to use the old Docker image tag that references the original backend resources. To fix this issue, the developer should update the Docker image tag to a new version and redeploy the application to the EKS cluster.
* A. The developer did not successfully create the new AWS account. This is incorrect. The creation of a new AWS account is not related to the application's connection to the
backend resources. The developer can use any AWS account to host the EKS cluster and the backend resources, as long as they have the proper permissions and configurations.
* B. The developer added a new tag to the Docker image. This is incorrect. Adding a new tag to the Docker image is not enough to deploy the changes to the application. The developer also needs to update the Docker image tag in the EKS cluster configuration, so that the EKS cluster can pull and run the new Docker image.
* D. The developer pushed the changes to a new Docker image tag. This is incorrect. Pushing the changes to a new Docker image tag is not enough to deploy the changes to the application. The developer also needs to update the Docker image tag in the EKS cluster configuration, so that the EKS cluster can pull and run the new Docker image. References:
? 1: Amazon EKS User Guide, "Deploying applications to your Amazon EKS
cluster", https://docs.aws.amazon.com/eks/latest/userguide/deploying- applications.html
? 2: Amazon ECR User Guide, "Pushing an image",
https://docs.aws.amazon.com/AmazonECR/latest/userguide/docker-push-ecr- image.html
? 3: Amazon EKS User Guide, "Updating an Amazon EKS cluster",
https://docs.aws.amazon.com/eks/latest/userguide/update-cluster.html

NEW QUESTION 144
A developer accesses AWS CodeCommit over SSH. The SSH keys configured to access AWS CodeCommit are tied to a user with the following permissions:

```
{
"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "codecommit:BatchGetRepositories",
      "codecommit:Get*",
      "codecommit:List*",
      "codecommit:GitPull"
    ],
    "Resource": "*"
  }
]
}
```

The developer needs to create/delete branches
Which specific IAM permissions need to be added based on the principle of least privilege?

A. "codecommit:CreateBranch"
   "codecommit:DeleteBranch"

B. "codecommit:Put*"

C. "codecommit:Update*"

D. "codecommit:*"

A. Option A
B. Option B
C. Option C
D. Option D

Answer: A

Explanation:
 This solution allows the developer to create and delete branches in AWS CodeCommit by granting the codecommit:CreateBranch and codecommit:DeleteBranch permissions. These are the minimum permissions required for this task, following the principle of least privilege. Option B grants too many permissions, such as codecommit:Put*, which allows the developer to create, update, or delete any resource in CodeCommit. Option C grants too few as permissions, such as codecommit:Update*, which does not allow the developer to create or delete branches. Option D grants all permissions, such as codecommit:*, which is not secure or recommended.
Reference: [AWS CodeCommit Permissions Reference], [Create a Branch (AWS CLI)]

NEW QUESTION 149
A developer needs to store configuration variables for an application. The developer needs to set an expiration date and time for me configuration. The developer wants to receive notifications. Before the configuration expires. Which solution will meet these requirements with the LEAST operational overhead?

A. Create a standard parameter in AWS Systems Manager Parameter Store Set Expiation and Expiration Notification policy types.
B. Create a standard parameter in AWS Systems Manager Parameter Store Create an AWS Lambda function to expire the configuration and to send Amazon Simple Notification Service (Amazon SNS) notifications.
C. Create an advanced parameter in AWS Systems Manager Parameter Store Set Expiration and Expiration Notification policy types.
D. Create an advanced parameter in AWS Systems Manager Parameter Store Create an Amazon EC2 instance with a corn job to expire the configuration and to send notifications.

Answer: C

Explanation:
 This solution will meet the requirements by creating an advanced parameter in AWS Systems Manager Parameter Store, which is a secure and scalable service for storing and managing configuration data and secrets. The advanced parameter allows setting expiration and expiration notification policy types, which enable specifying an expiration date and time for the configuration and receiving notifications before the configuration expires. The Lambda code will be refactored to load the Root CA Cert from the parameter store and modify the runtime trust store outside the Lambda function handler, which will improve performance and reduce latency by avoiding repeated calls to Parameter Store and trust store modifications for each invocation of the Lambda function. Option A is not optimal because it will create a standard parameter in AWS Systems Manager Parameter Store, which does not support expiration and expiration notification policy types. Option B is not optimal because it will create a secret access key and access key ID with permission to access the S3 bucket, which will introduce additional security risks and complexity for storing and managing credentials. Option D is not optimal because it will create a Docker container from Node.js base image to invoke Lambda functions, which will incur additional costs and overhead for creating and running Docker containers. References: AWS Systems Manager Parameter Store, [Using SSL/TLS to Encrypt a Connection to a DB Instance]

NEW QUESTION 151
A developer is troubleshooting an application mat uses Amazon DynamoDB in the uswest- 2 Region. The application is deployed to an Amazon EC2 instance. The application requires read-only permissions to a table that is named Cars The EC2 instance has an attached IAM role that contains the following IAM policy.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ReadOnlyAPIActions",
            "Effect": "Allow",
            "Action": [
                "dynamodb:GetItem",
                "dynamodb:BatchGetItem",
                "dynamodb:Scan",
                "dynamodb:Query",
                "dynamodb:ConditionCheckItem"
            ],
            "Resource": "arn:aws:dynamodb:us-west-2:account-id:table/Cars"
        }
    ]
}
```

When the application tries to read from the Cars table, an Access Denied error occurs. How can the developer resolve this error?

A. Modify the IAM policy resource to be "arn aws dynamo* us-west-2 account-id table/*"
B. Modify the IAM policy to include the dynamodb * action
C. Create a trust policy that specifies the EC2 service principa
D. Associate the role with the policy.
E. Create a trust relationship between the role and dynamodb Amazonas com.

Answer: C

Explanation:
 https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/access-control- overview.html#access-control-resource-ownership

NEW QUESTION 154
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DVA-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DVA-C02 Product From:

https://www.2passeasy.com/dumps/DVA-C02/

Money Back Guarantee

DVA-C02 Practice Exam Features:

* DVA-C02 Questions and Answers Updated Frequently

* DVA-C02 Practice Questions Verified by Expert Senior Certified Staff

* DVA-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DVA-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year