Amazon

Exam Questions DVA-C02

DVA-C02

NEW QUESTION 1
A developer is storing sensitive data generated by an application in Amazon S3. The developer wants to encrypt the data at rest. A company policy requires an audit trail of when the AWS Key Management Service (AWS KMS) key was used and by whom.
Which encryption option will meet these requirements?

A. Server-side encryption with Amazon S3 managed keys (SSE-S3)
Server-side encryption with AWS KMS managed keys (SSE-KMS}
B.
C. Server-side encryption with customer-provided keys (SSE-C)
D. Server-side encryption with self-managed keys

Answer: B

Explanation:
This solution meets the requirements because it encrypts data at rest using AWS KMS keys and provides an audit trail of when and by whom they were used. Server- side encryption with AWS KMS managed keys (SSE-KMS) is a feature of Amazon S3 that encrypts data using keys that are managed by AWS KMS. When SSE-KMS is enabled for an S3 bucket or object, S3 requests AWS KMS to generate data keys and encrypts data using these keys. AWS KMS logs every use of its keys in AWS CloudTrail, which records all API calls to AWS KMS as events. These events include information such as who made the request, when it was made, and which key was used. The company policy can use CloudTrail logs to audit critical events related to their data encryption and access. Server- side encryption with Amazon S3 managed keys (SSE-S3) also encrypts data at rest using keys that are managed by S3, but does not provide an audit trail of key usage. Server-side encryption with customer-provided keys (SSE-C) and server-side encryption with self- managed keys also encrypt data at rest using keys that are provided or managed by customers, but do not provide an audit trail of key usage and require additional overhead for key management.
Reference: [Protecting Data Using Server-Side Encryption with AWS KMS–Managed
Encryption Keys (SSE-KMS)], [Logging AWS KMS API calls with AWS CloudTrail]

NEW QUESTION 2
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
Modify the new API to add request validatio
C. Import the OpenAPI file.
D. Perform the test
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 3
A developer is incorporating AWS X-Ray into an application that handles personal
identifiable information (PII). The application is hosted on Amazon EC2 instances. The application trace messages include encrypted PII and go to Amazon CloudWatch. The developer needs to ensure that no PII goes outside of the EC2 instances.
Which solution will meet these requirements?

A. Manually instrument the X-Ray SDK in the application code.
B. Use the X-Ray auto-instrumentation agent.
C. Use Amazon Macie to detect and hide PI
D. Call the X-Ray API from AWS Lambda.
E. Use AWS Distro for Open Telemetry.

Answer: A

Explanation:
This solution will meet the requirements by allowing the developer to control what data is sent to X-Ray and CloudWatch from the application code. The developer can filter out any PII from the trace messages before sending them to X-Ray and CloudWatch, ensuring that no PII goes outside of the EC2 instances. Option B is

not optimal because it will automatically instrument all incoming and outgoing requests from the application, which may include PII in the trace messages. Option C is not optimal because it will require additional services and costs to use Amazon Macie and AWS Lambda, which may not be able to detect and hide all PII from the trace messages. Option D is not optimal because it will use Open Telemetry instead of X-Ray, which may not be compatible with CloudWatch and other AWS services.
References: [AWS X-Ray SDKs]

NEW QUESTION 4
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 5
A developer is creating an application that includes an Amazon API Gateway REST API in the us-east-2 Region. The developer wants to use Amazon CloudFront and a custom domain name for the API. The developer has acquired an SSL/TLS certificate for the domain from a third-party provider.
How should the developer configure the custom domain for the application?

A. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the same Region as the AP
B. Create a DNS A record for the custom domain.
C. Import the SSL/TLS certificate into CloudFron
D. Create a DNS CNAME record for the custom domain.
E. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the same Region as the AP
F. Create a DNS CNAME record for the custom domain.
G. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the us-east-1 Regio
H. Create a DNS CNAME record for the custom domain.

Answer: D

Explanation:
Amazon API Gateway is a service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. Amazon CloudFront is a content delivery network (CDN) service that can improve the performance and security of web applications. The developer can use CloudFront and a custom domain name for the API Gateway REST API. To do so, the developer needs to import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the us-east-1 Region. This is because CloudFront requires certificates from ACM to be in this Region. The developer also needs to create a DNS CNAME record for the custom domain that points to the CloudFront distribution.
References:
? [What Is Amazon API Gateway? - Amazon API Gateway]
? [What Is Amazon CloudFront? - Amazon CloudFront]
? [Custom Domain Names for APIs - Amazon API Gateway]

NEW QUESTION 6
A developer is deploying a new application to Amazon Elastic Container Service (Amazon ECS). The developer needs to securely store and retrieve different types of variables. These variables include authentication information for a remote API, the URL for the API, and credentials. The authentication information and API URL must be available to all current and future deployed versions of the application across development, testing, and production environments.
How should the developer retrieve the variables with the FEWEST application changes?

A. Update the application to retrieve the variables from AWS Systems Manager Parameter Stor
B. Use unique paths in Parameter Store for each variable in each environmen
C. Store the credentials in AWS Secrets Manager in each environment.
D. Update the application to retrieve the variables from AWS Key Management Service (AWS KMS). Store the API URL and credentials as unique keys for each environment.
E. Update the application to retrieve the variables from an encrypted file that is stored with the applicatio
F. Store the API URL and credentials in unique files for each environment.
G. Update the application to retrieve the variables from each of the deployed environment
H. Define the authentication information and API URL in the ECS task definition as unique names during the deployment process.

Answer: A

Explanation:
AWS Systems Manager Parameter Store is a service that provides secure, hierarchical storage for configuration data management and secrets management. The developer can update the application to retrieve the variables from Parameter Store by using the AWS SDK or the AWS CLI. The developer can use unique paths in Parameter Store for each variable in each environment, such as /dev/api-url, /test/api-url, and /prod/api-url. The developer can also store the credentials in AWS Secrets Manager, which is integrated with Parameter Store and provides additional features such as automatic rotation and encryption.
References:

? [What Is AWS Systems Manager? - AWS Systems Manager]
? [Parameter Store - AWS Systems Manager]
? [What Is AWS Secrets Manager? - AWS Secrets Manager]

NEW QUESTION 7
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 8
A developer is testing a RESTful application that is deployed by using Amazon API Gateway and AWS Lambda When the developer tests the user login by using credentials that are not valid, the developer receives an HTTP 405 METHOD_NOT_ALLOWED error The developer has verified that the test is sending the correct request for the resource
Which HTTP error should the application return in response to the request?

A. HTTP 401
B. HTTP 404
C. HTTP 503
D. HTTP 505

Answer: A

Explanation:
The HTTP 401 error indicates that the request has not been applied because it lacks valid authentication credentials for the target resource. This is the appropriate error code to return when the user login fails due to invalid credentials. The HTTP 405 error means that the method specified in the request is not allowed for the resource identified by the request URI, which is not the case here. The other error codes are not relevant to the authentication failure scenario.
References
? HTTP Status Codes
? AWS Lambda Function Errors in API Gateway

NEW QUESTION 9
A company is building a new application that runs on AWS and uses Amazon API Gateway to expose APIs Teams of developers are working on separate components of the application in parallel.The company wants to publish an API without an integrated backend so that teams that depend on the application backend can continue the development work before the API backend development is complete.
Which solution will meet these requirements?

A. Create API Gateway resources and set the integration type value to MOCK Configure the method integration request and integration response to associate a response with an HTTP status code Create an API Gateway stage and deploy the API.
B. Create an AWS Lambda function that returns mocked responses and various HTTP status code
C. Create API Gateway resources and set the integration type value to AWS_PROXY Deploy the API.
D. Create an EC2 application that returns mocked HTTP responses Create API Gateway resources and set the integration type value to AWS Create an API Gateway stage and deploy the API.
E. Create API Gateway resources and set the integration type value set to HTTP_PROX
F. Add mapping templates and deploy the AP
G. Create an AWS Lambda layer that returns various HTTP status codes Associate the Lambda layer with the API deployment

Answer: A

Explanation:
The best solution for publishing an API without an integrated backend is to use the MOCK integration type in API Gateway. This allows the developer to return a static response to the client without sending the request to a backend service. The developer can configure the method integration request and integration response to associate a response with an HTTP status code, such as 200 OK or 404 Not Found. The developer can also create an API Gateway stage and deploy the API to make it available to the teams that depend on the application backend. The other solutions are either not feasible or not efficient. Creating an AWS Lambda function, an EC2 application, or an AWS Lambda layer would require additional resources and code to generate the mocked responses and HTTP status codes. These solutions would also incur additional costs and complexity, and would not leverage the built-in functionality of API Gateway. References
? Set up mock integrations for API Gateway REST APIs
? Mock Integration for API Gateway - AWS CloudFormation
? Mocking API Responses with API Gateway
? How to mock API Gateway responses with AWS SAM

NEW QUESTION 10
A mobile app stores blog posts in an Amazon DynacnoDB table Millions of posts are added every day and each post represents a single item in the table. The

mobile app requires only recent posts. Any post that is older than 48 hours can be removed.
What is the MOST cost-effective way to delete posts that are older man 48 hours?

A. For each item add a new attribute of type String that has a timestamp that is set to the blog post creation tim
B. Create a script to find old posts with a table scan and remove posts that are order than 48 hours by using the Balch Write Item API operatio
C. Schedule a cron job on an Amazon EC2 instance once an hour to start the script.
D. For each item add a new attribute of typ
E. String that has a timestamp that its set to the blog post creation tim
F. Create a script to find old posts with a table scan and remove posts that are Oder than 48 hours by using the Batch Write item API operatin
G. Place the script in a container imag
H. Schedule an Amazon Elastic Container Service (Amazon ECS) task on AWS Far gate that invokes the container every 5 minutes.
I. For each item, add a new attribute of type Date that has a timestamp that is set to 48 hours after the blog post creation tim
J. Create a global secondary index (GSI) that uses the new attribute as a sort ke
K. Create an AWS Lambda function that references the GSI and removes expired items by using the Batch Write item API operation Schedule me function with an Amazon CloudWatch event every minute.
L. For each item add a new attribute of typ
M. Number that has timestamp that is set to 48 hours after the blog pos
N. creation time Configure the DynamoDB table with a TTL that references the new attribute.

Answer: D

Explanation:
This solution will meet the requirements by using the Time to Live (TTL) feature of DynamoDB, which enables automatically deleting items from a table after a certain time period. The developer can add a new attribute of type Number that has a timestamp that is set to 48 hours after the blog post creation time, which represents the expiration time of the item. The developer can configure the DynamoDB table with a TTL that references the new attribute, which instructs DynamoDB to delete the item when the current time is greater than or equal to the expiration time. This solution is also cost- effective as it does not incur any additional charges for deleting expired items. Option A is not optimal because it will create a script to find and remove old posts with a table scan and a batch write item API operation, which may consume more read and write capacity units and incur more costs. Option B is not optimal because it will use Amazon Elastic Container Service (Amazon ECS) and AWS Fargate to run the script, which may introduce additional costs and complexity for managing and scaling containers. Option C is not optimal because it will create a global secondary index (GSI) that uses the expiration time as a sort key, which may consume more storage space and incur more costs.
References: Time To Live, Managing DynamoDB Time To Live (TTL)

NEW QUESTION 10
A company is offering APIs as a service over the internet to provide unauthenticated read access to statistical information that is updated daily. The company uses Amazon API Gateway and AWS Lambda to develop the APIs. The service has become popular, and the company wants to enhance the responsiveness of the APIs.
Which action can help the company achieve this goal?

A. Enable API caching in API Gateway.
B. Configure API Gateway to use an interface VPC endpoint.
C. Enable cross-origin resource sharing (CORS) for the APIs.
D. Configure usage plans and API keys in API Gateway.

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 enable API caching in API Gateway to cache responses from the backend integration point for a specified time-to- live (TTL) period. This can improve the responsiveness of the APIs by reducing the number of calls made to the backend service. References:
? [What Is Amazon API Gateway? - Amazon API Gateway]
? [Enable API Caching to Enhance Responsiveness - Amazon API Gateway]

NEW QUESTION 11
A developer is optimizing an AWS Lambda function and wants to test the changes in production on a small percentage of all traffic. The Lambda function serves requests to a REST API in Amazon API Gateway. The developer needs to deploy their changes and perform a test in production without changing the API Gateway URL.
Which solution will meet these requirements?

A. Define a function version for the currently deployed production Lambda functio
B. Update the API Gateway endpoint to reference the new Lambda function versio
C. Upload and publish the optimized Lambda function cod
D. On the production API Gateway stage, define a canary release and set the percentage of traffic to direct to the canary releas
E. Update the API Gateway endpoint to use the $LATEST version of the Lambda functio
F. Publish the API to the canary stage.
G. Define a function version for the currently deployed production Lambda functio
H. Update the API Gateway endpoint to reference the new Lambda function versio
I. Upload and publish the optimized Lambda function cod
J. Update the API Gateway endpoint to use the $LATEST version of the Lambda functio
K. Deploy a new API Gateway stage.
L. Define an alias on the $LATEST version of the Lambda functio
M. Update the API Gateway endpoint to reference the new Lambda function alia
N. Upload and publish the optimized Lambda function cod
O. On the production API Gateway stage, define a canary release and set the percentage of traffic to direct to the canary releas
P. Update the API Gateway endpoint to use the SLAT EST version of the Lambda functio
Q. Publish to the canary stage.
R. Define a function version for the currently deployed production Lambda functio
S. Update the API Gateway endpoint to reference the new Lambda function versio
T. Upload and publish the optimized Lambda function cod
. Update the API Gateway endpoint to use the$LATEST version of the Lambda functio

. Deploy the API to the production API Gateway stage.

Answer: C

Explanation:
? A Lambda alias is a pointer to a specific Lambda function version or another alias1. A Lambda alias allows you to invoke different versions of a function using the same name1. You can also split traffic between two aliases by assigning weights to them1.
? In this scenario, the developer needs to test their changes in production on a small percentage of all traffic without changing the API Gateway URL. To achieve this, the developer can follow these steps:
? By using this solution, the developer can test their changes in production on a small percentage of all traffic without changing the API Gateway URL. The developer can also monitor and compare metrics between the canary and production releases, and promote or disable the canary as needed2.

NEW QUESTION 15
A developer needs to build an AWS CloudFormation template that self-populates the AWS Region variable that deploys the CloudFormation template
What is the MOST operationally efficient way to determine the Region in which the template is being deployed?

A. Use the AWS:.Region pseudo parameter
B. Require the Region as a CloudFormation parameter
C. Find the Region from the AWS::StackId pseudo parameter by using the Fn::Split intrinsic function
D. Dynamically import the Region by referencing the relevant parameter in AWS Systems Manager Parameter Store

Answer: A

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/mappings-section- structure.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/pseudo-parameter- reference.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/pseudo-parameter- reference.html

NEW QUESTION 20
An Amazon Kinesis Data Firehose delivery stream is receiving customer data that contains personally identifiable information. A developer needs to remove pattern-based customer identifiers from the data and store the modified data in an Amazon S3 bucket.
What should the developer do to meet these requirements?

A. Implement Kinesis Data Firehose data transformation as an AWS Lambda functio
B. Configure the function to remove the customer identifier
C. Set an Amazon S3 bucket as the destination of the delivery stream.
D. Launch an Amazon EC2 instanc
E. Set the EC2 instance as the destination of the delivery strea
F. Run an application on the EC2 instance to remove the customer identifier
G. Store the transformed data in an Amazon S3 bucket.
H. Create an Amazon OpenSearch Service instanc
I. Set the OpenSearch Service instance as the destination of the delivery strea
J. Use search and replace to remove the customer identifier
K. Export the data to an Amazon S3 bucket.
L. Create an AWS Step Functions workflow to remove the customer identifier
M. As the last step in the workflow, store the transformed data in an Amazon S3 bucke
N. Set the workflow as the destination of the delivery stream.

Answer: A

Explanation:
Amazon Kinesis Data Firehose is a service that delivers real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon OpenSearch Service, and Amazon Kinesis Data Analytics. The developer can implement Kinesis Data Firehose data transformation as an AWS Lambda function. The function can remove pattern-based customer identifiers from the data and return the modified data to Kinesis Data Firehose. The developer can set an Amazon S3 bucket as the destination of the delivery stream. References:
? [What Is Amazon Kinesis Data Firehose? - Amazon Kinesis Data Firehose]
? [Data Transformation - Amazon Kinesis Data Firehose]

NEW QUESTION 22
A company notices that credentials that the company uses to connect to an external software as a service (SaaS) vendor are stored in a configuration file as plaintext.
The developer needs to secure the API credentials and enforce automatic credentials rotation on a quarterly basis.
Which solution will meet these requirements MOST securely?

A. Use AWS Key Management Service (AWS KMS) to encrypt the configuration fil
B. Decrypt the configuration file when users make API calls to the SaaS vendo
C. Enable rotation.
D. Retrieve temporary credentials from AWS Security Token Service (AWS STS) every 15 minute
E. Use the temporary credentials when users make API calls to the SaaS vendor.
F. Store the credentials in AWS Secrets Manager and enable rotatio
G. Configure the API to have Secrets Manager access.
H. Store the credentials in AWS Systems Manager Parameter Store and enable rotatio
I. Retrieve the credentials when users make API calls to the SaaS vendor.

Answer: C

Explanation:
Store the credentials in AWS Secrets Manager and enable rotation. Configure the API to have Secrets Manager access. This is correct. This solution will meet the

requirements most securely, because it uses a service that is designed to store and manage secrets such as API credentials. AWS Secrets Manager helps you protect access to your applications, services, and IT resources by enabling you to rotate, manage, and retrieve secrets throughout their lifecycle1. You can store secrets such as passwords, database strings, API keys, and license codes as encrypted values2. You can also configure automatic rotation of your secrets on a schedule that you specify3. You can use the AWS SDK or CLI to retrieve secrets from Secrets Manager when you need them4. This way, you can avoid storing credentials in plaintext files or hardcoding them in your code.

NEW QUESTION 27
A company has a multi-node Windows legacy application that runs on premises. The application uses a network shared folder as a centralized configuration repository to store configuration files in .xml format. The company is migrating the application to Amazon EC2 instances. As part of the migration to AWS, a developer must identify a solution that provides high availability for the repository.
Which solution will meet this requirement MOST cost-effectively?

A. Mount an Amazon Elastic Block Store (Amazon EBS) volume onto one of the EC2 instance
B. Deploy a file system on the EBS volum
C. Use the host operating system to share a folde
D. Update the application code to read and write configuration files from the shared folder.
E. Deploy a micro EC2 instance with an instance store volum
F. Use the host operating system to share a folde
G. Update the application code to read and write configuration files from the shared folder.
H. Create an Amazon S3 bucket to host the repositor
I. Migrate the existing .xml files to the S3 bucke
J. Update the application code to use the AWS SDK to read and write configuration files from Amazon S3.
K. Create an Amazon S3 bucket to host the repositor
L. Migrate the existing .xml files to the S3 bucke
M. Mount the S3 bucket to the EC2 instances as a local volum
N. Update the application code to read and write configuration files from the disk.

Answer: C

Explanation:
Amazon S3 is a service that provides highly scalable, durable, and secure object storage. The developer can create an S3 bucket to host the repository and migrate the existing .xml files to the S3 bucket. The developer can update the application code to use the AWS SDK to read and write configuration files from S3. This solution will meet the requirement of high availability for the repository in a cost-effective way.
References:
? [Amazon Simple Storage Service (S3)]
? [Using AWS SDKs with Amazon S3]

NEW QUESTION 31
A developer has been asked to create an AWS Lambda function that is invoked any time updates are made to items in an Amazon DynamoDB table. The function has been created and appropriate permissions have been added to the Lambda execution role Amazon DynamoDB streams have been enabled for the table, but the function 15 still not being invoked.
Which option would enable DynamoDB table updates to invoke the Lambda function?

A. Change the StreamViewType parameter value to NEW_AND_OLOJMAGES for the DynamoDB table.
B. Configure event source mapping for the Lambda function.
C. Map an Amazon Simple Notification Service (Amazon SNS) topic to the DynamoDB streams.
D. Increase the maximum runtime (timeout) setting of the Lambda function.

Answer: B

Explanation:
This solution allows the Lambda function to be invoked by the DynamoDB stream whenever updates are made to items in the DynamoDB table. Event source mapping is a feature of Lambda that enables a function to be triggered by an event source, such as a DynamoDB stream, an Amazon Kinesis stream, or an Amazon Simple Queue Service (SQS) queue. The developer can configure event source mapping for the Lambda function using the AWS Management Console, the AWS CLI, or the AWS SDKs. Changing the StreamViewType parameter value to NEW_AND_OLD_IMAGES for the DynamoDB table will not affect the invocation of the Lambda function, but only change the information that is written to the stream record. Mapping an Amazon Simple Notification Service (Amazon SNS) topic to the DynamoDB stream will not invoke the Lambda function directly, but require an additional subscription from the Lambda function to the SNS topic. Increasing the maximum runtime (timeout) setting of the Lambda function will not affect the invocation of the Lambda function, but only change how long the function can run before it is terminated.
Reference: [Using AWS Lambda with Amazon DynamoDB], [Using AWS Lambda with Amazon SNS]

NEW QUESTION 36
An application that runs on AWS receives messages from an Amazon Simple Queue Service (Amazon SQS) queue and processes the messages in batches. The application sends the data to another SQS queue to be consumed by another legacy application. The legacy system can take up to 5 minutes to process some transaction data.
A developer wants to ensure that there are no out-of-order updates in the legacy system. The developer cannot alter the behavior of the legacy system.
Which solution will meet these requirements?

A. Use an SQS FIFO queu
B. Configure the visibility timeout value.
C. Use an SQS standard queue with a SendMessageBatchRequestEntry data typ
D. Configure the DelaySeconds values.
E. Use an SQS standard queue with a SendMessageBatchRequestEntry data typ
F. Configure the visibility timeout value.
G. Use an SQS FIFO queu
H. Configure the DelaySeconds value.

Answer: A

Explanation:

? An SQS FIFO queue is a type of queue that preserves the order of messages and ensures that each message is delivered and processed only once1. This is suitable for the scenario where the developer wants to ensure that there are no out-of-order updates in the legacy system.
? The visibility timeout value is the amount of time that a message is invisible in the queue after a consumer receives it2. This prevents other consumers from processing the same message simultaneously. If the consumer does not delete the message before the visibility timeout expires, the message becomes visible again and another consumer can receive it2.
? In this scenario, the developer needs to configure the visibility timeout value to be longer than the maximum processing time of the legacy system, which is 5 minutes. This will ensure that the message remains invisible in the queue until the legacy system finishes processing it and deletes it. This will prevent duplicate or out-of-order processing of messages by the legacy system.

NEW QUESTION 37
A developer is creating an AWS Lambda function that needs credentials to connect to an Amazon RDS for MySQL database. An Amazon S3 bucket currently stores the credentials. The developer needs to improve the existing solution by implementing credential rotation and secure storage. The developer also needs to provide integration with the Lambda function.
Which solution should the developer use to store and retrieve the credentials with the LEAST management overhead?

A. Store the credentials in AWS Systems Manager Parameter Stor
B. Select the database that the parameter will acces
C. Use the default AWS Key Management Service (AWS KMS) key to encrypt the paramete
D. Enable automatic rotation for the paramete
E. Use the parameter from Parameter Store on the Lambda function to connect to the database.
F. Encrypt the credentials with the default AWS Key Management Service (AWS KMS) ke
G. Store the credentials as environment variables for the Lambda functio
H. Create a second Lambda function to generate new credentials and to rotate the credentials by updating the environment variables of the first Lambda functio
I. Invoke the second Lambda function by using an Amazon EventBridge rule that runs on a schedul
J. Update the database to use the new credential
K. On the first Lambda function, retrieve the credentials from the environment variable
L. Decrypt the credentials by using AWS KMS, Connect to the database.
M. Store the credentials in AWS Secrets Manage
N. Set the secret type to Credentials for Amazon RDS databas
O. Select the database that the secret will acces
P. Use the default AWS Key Management Service (AWS KMS) key to encrypt the secre
Q. Enable automatic rotation for the secre
R. Use the secret from Secrets Manager on the Lambda function to connect to the database.
S. Encrypt the credentials by using AWS Key Management Service (AWS KMS). Store the credentials in an Amazon DynamoDB tabl
T. Create a second Lambda function to rotate the credential
. Invoke the second Lambda function by using an Amazon EventBridge rule that runs on a schedul
. Update the DynamoDB tabl
. Update the database to use the generated credential
. Retrieve the credentials from DynamoDB with the first Lambda functio
. Connect to the database.

Answer: C

Explanation:
AWS Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources. Secrets Manager enables you to store, retrieve, and rotate secrets such as database credentials, API keys, and passwords. Secrets Manager supports a secret type for RDS databases, which allows you to select an existing RDS database instance and generate credentials for it. Secrets Manager encrypts the secret using AWS Key Management Service (AWS KMS) keys and enables automatic rotation of the secret at a specified interval. A Lambda function can use the AWS SDK or CLI to retrieve the secret from Secrets Manager and use it to connect to the database. Reference: Rotating your AWS Secrets Manager secrets

NEW QUESTION 39
A developer has observed an increase in bugs in the AWS Lambda functions that a development team has deployed in its Node is application. To minimize these bugs, the developer wants to impendent automated testing of Lambda functions in an environment that Closely simulates the Lambda environment.
The developer needs to give other developers the ability to run the tests locally. The developer also needs to integrate the tests into the team's continuous integration and continuous delivery (Ct/CO) pipeline before the AWS Cloud Development Kit (AWS COK) deployment.
Which solution will meet these requirements?

A. Create sample events based on the Lambda documentatio
B. Create automated test scripts that use the cdk local invoke command to invoke the Lambda function
C. Check the response Document the test scripts for the other developers on the team Update the CI/CD pipeline to run the test scripts.
D. Install a unit testing framework that reproduces the Lambda execution environmen
E. Create sample events based on the Lambda Documentation Invoke the handler function by using a unit testing framewor
framework for the other developers on the tea
F. Check the response Document how to run the unit testing
G. Update the CCD pipeline to run the unit testing framework.
H. Install the AWS Serverless Application Model (AWS SAW) CLI tool Use the Sam local generate-event command to generate sample events for me automated test
I. Create automated test scripts that use the Sam local invoke command to invoke the Lambda function
J. Check the response Document the test scripts tor the other developers on the team Update the CI/CD pipeline to run the test scripts.
K. Create sample events based on the Lambda documentatio
L. Create a Docker container from the Node is base image to invoke the Lambda function
M. Check the response Document how to run the Docker container for the more developers on the team update the CI/CD pipeline to run the Docker container.

Answer: C

Explanation:
This solution will meet the requirements by using AWS SAM CLI tool, which is a command line tool that lets developers locally build, test, debug, and deploy serverless applications defined by AWS SAM templates. The developer can use sam local generate- event command to generate sample events for different event sources such as API Gateway or S3. The developer can create automated test scripts that use sam local invoke command to invoke Lambda functions locally in an environment that closely simulates Lambda environment. The developer can check the response from Lambda functions and document how to run the test scripts for other developers on the team. The developer can also update CI/CD pipeline to run these test scripts before deploying with AWS CDK. Option A is not

optimal because it will use cdk local invoke command, which does not exist in AWS CDK CLI tool. Option B is not optimal because it will use a unit testing framework that reproduces Lambda execution environment, which may not be accurate or consistent with Lambda environment. Option D is not optimal because it will create a Docker container from Node.js base image to invoke Lambda functions, which may introduce additional overhead and complexity for creating and running Docker containers.
References: [AWS Serverless Application Model (AWS SAM)], [AWS Cloud Development Kit (AWS CDK)]

NEW QUESTION 44
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 46
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 50
An online sales company is developing a serverless application that runs on AWS. The application uses an AWS Lambda function that calculates order success rates and stores the data in an Amazon DynamoDB table. A developer wants an efficient way to invoke the Lambda function every 15 minutes.
Which solution will meet this requirement with the LEAST development effort?

A. Create an Amazon EventBridge rule that has a rate expression that will run the rule every 15 minute
B. Add the Lambda function as the target of the EventBridge rule.
C. Create an AWS Systems Manager document that has a script that will invoke the Lambda function on Amazon EC2. Use a Systems Manager Run Command task to run the shell script every 15 minutes.
D. Create an AWS Step Functions state machin
E. Configure the state machine to invoke the Lambda function execution role at a specified interval by using a Wait stat
F. Set the interval to 15 minutes.
G. Provision a small Amazon EC2 instanc
H. Set up a cron job that invokes the Lambda function every 15 minutes.

Answer: A

Explanation:
The best solution for this requirement is option A. Creating an Amazon EventBridge rule that has a rate expression that will run the rule every 15 minutes and adding the Lambda function as the target of the EventBridge rule is the most efficient way to invoke the Lambda function periodically. This solution does not require any additional resources or development effort, and it leverages the built-in scheduling capabilities of EventBridge1.

NEW QUESTION 53
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 55
A developer designed an application on an Amazon EC2 instance The application makes API requests to objects in an Amazon S3 bucket
Which combination of steps will ensure that the application makes the API requests in the MOST secure manner? (Select TWO.)

A. Create an IAM user that has permissions to the S3 bucke
B. Add the user to an 1AM group
C. Create an IAM role that has permissions to the S3 bucket
D. Add the IAM role to an instance profil
E. Attach the instance profile to the EC2 instance.
F. Create an 1AM role that has permissions to the S3 bucket Assign the role to an 1AM group
G. Store the credentials of the IAM user in the environment variables on the EC2 instance

Answer: BC

Explanation:
- Create an IAM role that has permissions to the S3 bucket. - Add the IAM role to an instance profile. Attach the instance profile to the EC2 instance. We first need to create a n IAM Role with permissions to read and eventually write a specific S3 bucket. Then, we need to attach the role to the EC2 isntance through an instance profile. In this
way, the ec2 instance has the permissions to read and eventually write the specified S3 bucket

NEW QUESTION 57
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 61
A developer has written the following IAM policy to provide access to an Amazon S3 bucket:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "s3:GetObject",
                "s3:PutObject"
            ],
            "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/secrets*"
        }
    ]
}
```

Which access does the policy allow regarding the s3:GetObject and s3:PutObject actions?

A. Access on all buckets except the “DOC-EXAMPLE-BUCKET” bucket
B. Access on all buckets that start with “DOC-EXAMPLE-BUCKET” except the “DOC-EXAMPLE-BUCKET/secrets” bucket
C. Access on all objects in the “DOC-EXAMPLE-BUCKET” bucket along with access to all S3 actions for objects in the “DOC-EXAMPLE-BUCKET” bucket that start with “secrets”
D. Access on all objects in the “DOC-EXAMPLE-BUCKET” bucket except on objects that start with “secrets”

Answer: D

Explanation:
The IAM policy shown in the image is a resource-based policy that grants or denies access to an S3 bucket based on certain conditions. The first statement allows access to any S3 action on any object in the “DOC-EXAMPLE-BUCKET” bucket when the request is made over HTTPS (the value of aws:SecureTransport is true). The second statement denies access to the s3:GetObject and s3:PutObject actions on any object in the “DOC-EXAMPLE-BUCKET/secrets” prefix when the request is made over HTTP (the value of aws:SecureTransport is false). Therefore, the policy allows access on all objects in the “DOC-EXAMPLE-BUCKET” bucket except on objects that start with “secrets”.
Reference: Using IAM policies for Amazon S3

NEW QUESTION 64
A company receives food orders from multiple partners. The company has a microservices application that uses Amazon API Gateway APIs with AWS Lambda integration. Each partner sends orders by calling a customized API that is exposed through API Gateway. The API call invokes a shared Lambda function to process the orders.
Partners need to be notified after the Lambda function processes the orders. Each partner must receive updates for only the partner's own orders. The company wants to add new partners in the future with the fewest code changes possible.
Which solution will meet these requirements in the MOST scalable way?

A. Create a different Amazon Simple Notification Service (Amazon SNS) topic for each partne
B. Configure the Lambda function to publish messages for each partner to the partner's SNS topic.
C. Create a different Lambda function for each partne
D. Configure the Lambda function to notify each partner's service endpoint directly.
E. Create an Amazon Simple Notification Service (Amazon SNS) topi
F. Configure the Lambda function to publish messages with specific attributes to the SNS topi
G. Subscribe each partner to the SNS topi
H. Apply the appropriate filter policy to the topic subscriptions.
Create one Amazon Simple Notification Service (Amazon SNS) topi
I. Subscribe all partners to the SNS topic.

Answer: C

Explanation:
Amazon Simple Notification Service (Amazon SNS) is a fully managed messaging service that enables pub/sub communication between distributed systems. The developer can create an SNS topic and configure the Lambda function to publish messages with specific attributes to the topic. The developer can subscribe each partner to the SNS topic and apply the appropriate filter policy to the topic subscriptions. This way, each partner will receive updates for only their own orders based on the message attributes. This solution will meet the requirements in the most scalable way and allow adding new partners in the future with minimal code changes.
References:
? [Amazon Simple Notification Service (SNS)]
? [Filtering Messages with Attributes - Amazon Simple Notification Service]

NEW QUESTION 65
A developer is creating an application that will give users the ability to store photos from their cellphones in the cloud. The application needs to support tens of thousands of users. The application uses an Amazon API Gateway REST API that is integrated with AWS Lambda functions to process the photos. The application stores details about the photos in Amazon DynamoDB.
Users need to create an account to access the application. In the application, users must be able to upload photos and retrieve previously uploaded photos. The photos will range in size from 300 KB to 5 MB.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon Cognito user pools to manage user account
B. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the AP
C. Use the Lambda function to store the photos and details in the DynamoDB tabl
D. Retrieve previously uploaded photos directly from the DynamoDB table.
E. Use Amazon Cognito user pools to manage user account
F. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the AP
G. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the DynamoDB tabl
H. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
I. Create an IAM user for each user of the application during the sign-up proces
J. Use IAM authentication to access the API Gateway AP
DynamoDB
K. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the tabl
L. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
M. Create a users table in DynamoD
N. Use the table to manage user account
O. Create a Lambda authorizer that validates user credentials against the users tabl
P. Integrate the Lambda authorizer with API Gateway to control access to the AP
Q. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as par of the photo details in the DynamoDB tabl
R. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.

Answer: B

Explanation:
Amazon Cognito user pools is a service that provides a secure user directory that scales to hundreds of millions of users. The developer can use Amazon Cognito user pools to manage user accounts and create an Amazon Cognito user pool authorizer in API Gateway to control access to the API. The developer can use the Lambda function to store the photos in Amazon S3, which is a highly scalable, durable, and secure object storage service. The developer can store the object’s S3 key as part of the photo details in the DynamoDB table, which is a fast and flexible NoSQL database service. The developer can retrieve previously uploaded photos by querying DynamoDB for the S3 key and fetching the photos from S3. This solution will meet the requirements with the least operational overhead.
References:
? [Amazon Cognito User Pools]
? [Use Amazon Cognito User Pools - Amazon API Gateway]
? [Amazon Simple Storage Service (S3)]

? [Amazon DynamoDB]

NEW QUESTION 68
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 70
A developer must use multi-factor authentication (MFA) to access data in an Amazon S3 bucket that is in another AWS account. Which AWS Security Token Service (AWS STS) API operation should the developer use with the MFA information to meet this requirement?

A. AssumeRoleWithWebidentity
B. GetFederationToken
C. AssumeRoleWithSAML
D. AssumeRole

Answer: D

Explanation:
The AssumeRole API operation returns a set of temporary security credentials that can be used to access resources in another AWS account. The developer can specify the MFA device serial number and the MFA token code in the request parameters. This option enables the developer to use MFA to access data in an S3 bucket that is in another AWS account. The other options are not relevant or effective for this scenario. References
? AssumeRole
? Requesting Temporary Security Credentials

NEW QUESTION 75
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 AWS Serverless Application Model Command Line Interface (AWS SAM CLI) is a command-line tool for local development and testing of Serverless applications2. The sam local start-api subcommand of AWS SAM CLI is used to simulate a REST API by starting a new local endpoint3. Therefore, option D is correct.

NEW QUESTION 78
A developer is trying get data from an Amazon DynamoDB table called demoman-table. The developer configured the AWS CLI to use a specific IAM use's credentials and ran the following command.

```
aws dynamodb get-item --table-name demoman-table --key '{"id": {"N":"1993"}}'
```

The command returned errors and no rows were returned. What is the MOST likely cause of these issues?

A. The command is incorrect; it should be rewritten to use put-item with a string argument
B. The developer needs to log a ticket with AWS Support to enable access to the demoman-table
C. Amazon DynamoOB cannot be accessed from the AWS CLI and needs to called via the REST API
D. The IAM user needs an associated policy with read access to demoman-table

Answer: D

Explanation:
This solution will most likely solve the issues because it will grant the IAM user the necessary permission to access the DynamoDB table using the AWS CLI command. The error message indicates that the IAM user does not have sufficient access rights to perform the scan operation on the table. Option A is not optimal because it will change the command to use put-item instead of scan, which will not achieve the desired result of getting data from the table. Option B is not optimal because it will involve contacting AWS Support, which may not be necessary or efficient for this issue. Option C is not optimal because it will state that DynamoDB cannot be accessed from the AWS CLI, which is incorrect as DynamoDB supports AWS CLI commands.
References: AWS CLI for DynamoDB, [IAM Policies for DynamoDB]

NEW QUESTION 81
A company built an online event platform For each event the company organizes quizzes and generates leaderboards that are based on the quiz scores. The company stores the leaderboard data in Amazon DynamoDB and retains the data for 30 days after an event is complete The company then uses a scheduled job to delete the old leaderboard data
The DynamoDB table is configured with a fixed write capacity. During the months when many events occur, the DynamoDB write API requests are throttled when the scheduled delete job runs.
A developer must create a long-term solution that deletes the old leaderboard data and optimizes write throughput
Which solution meets these requirements?

A. Configure a TTL attribute for the leaderboard data
B. Use DynamoDB Streams to schedule and delete the leaderboard data
C. Use AWS Step Functions to schedule and delete the leaderboard data.
D. Set a higher write capacity when the scheduled delete job runs

Answer: A

Explanation:
"deletes the item from your table without consuming any write throughput" https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html

NEW QUESTION 84
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 89
A company is running a custom application on a set of on-premises Linux servers that are accessed using Amazon API Gateway. AWS X-Ray tracing has been enabled on the API test stage.
How can a developer enable X-Ray tracing on the on-premises servers with the LEAST amount of configuration?

A. Install and run the X-Ray SDK on the on-premises servers to capture and relay the data to the X-Ray service.
B. Install and run the X-Ray daemon on the on-premises servers to capture and relay the data to the X-Ray service.
C. Capture incoming requests on-premises and configure an AWS Lambda function to pull, process, and relay relevant data to X-Ray using the PutTraceSegments API call.
D. Capture incoming requests on-premises and configure an AWS Lambda function to pull, process, and relay relevant data to X-Ray using the PutTelemetryRecords API call.

Answer: B

Explanation:
The X-Ray daemon is a software that collects trace data from the X-Ray SDK and relays it to the X-Ray service. The X-Ray daemon can run on any platform that supports Go, including Linux, Windows, and macOS. The developer can install and run the X-Ray daemon on the on-premises servers to capture and relay the data to the X-Ray service with minimal configuration. The X-Ray SDK is used to instrument the application code, not to capture and relay data. The Lambda function solutions are more complex and require additional configuration.
References:
? [AWS X-Ray concepts - AWS X-Ray]
? [Setting up AWS X-Ray - AWS X-Ray]

NEW QUESTION 94

A developer warns 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 lest 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 API Import the OpenAPI file Modify the new API to add request validatio
C. Perform the tests Modify the existing API to add request validatio
D. Deploy the existing API to production.
E. Modify the existing API to add request validatio
F. Deploy the updated API to a new API Gateway stage Perform the tests Deploy the updated API to the API Gateway production stage.
G. Create a new API Add the necessary resources and methods including new request validatio
H. Perform the tests Modify the existing API to add request validatio
I. Deploy the existing API to production.
J. Clone the exiting API Modify the new API lo add request validatio
Modify the existing API to add request validation Deploy the existing API to production.
K. Perform the tests

Answer: D

Explanation:
This solution allows the developer to test the changes without affecting the production environment. Cloning an API creates a copy of the API definition that can be modified independently. The developer can then add request validation to the new API and test it using a testing tool. After verifying that the changes work as expected, the developer can apply the same changes to the existing API and deploy it to production.
Reference: Clone an API, [Enable Request Validation for an API in API Gateway]

NEW QUESTION 99
A developer is planning to migrate on-premises company data to Amazon S3. The data must be encrypted, and the encryption Keys must support automate annual rotation. The company must use AWS Key Management Service (AWS KMS) to encrypt the data.
When type of keys should the developer use to meet these requirements?

A. Amazon S3 managed keys
B. Symmetric customer managed keys with key material that is generated by AWS
C. Asymmetric customer managed keys with key material that generated by AWS
D. Symmetric customer managed keys with imported key material

Answer: B

Explanation:
The type of keys that the developer should use to meet the requirements is symmetric customer managed keys with key material that is generated by AWS. This way, the developer can use AWS Key Management Service (AWS KMS) to encrypt the data with a symmetric key that is managed by the developer. The developer can also enable automatic annual rotation for the key, which creates new key material for the key every year. The other options either involve using Amazon S3 managed keys, which do not support automatic annual rotation, or using asymmetric keys or imported key material, which are not supported by S3 encryption.
Reference: Using AWS KMS keys to encrypt S3 objects

NEW QUESTION 102
A developer is creating an Amazon DynamoDB table by using the AWS CLI The DynamoDB table must use server-side encryption with an AWS owned encryption key
How should the developer create the DynamoDB table to meet these requirements?

A. Create an AWS Key Management Service (AWS KMS) customer managed ke
B. Provide the key's Amazon Resource Name (ARN) in the KMSMasterKeyld parameter during creation of the DynamoDB table
C. Create an AWS Key Management Service (AWS KMS) AWS managed key Provide the key's Amazon Resource Name (ARN) in the KMSMasterKeyld parameter during creation of the DynamoDB table
D. Create an AWS owned key Provide the key's Amazon Resource Name (ARN) in the KMSMasterKeyld parameter during creation of the DynamoDB table.
E. Create the DynamoDB table with the default encryption options

Answer: D

Explanation:
When creating an Amazon DynamoDB table using the AWS CLI, server-side encryption with an AWS owned encryption key is enabled by default. Therefore, the developer does not need to create an AWS KMS key or specify the KMSMasterKeyId parameter. Option A and B are incorrect because they suggest creating customer- managed and AWS-managed KMS keys, which are not needed in this scenario. Option C is also incorrect because AWS owned keys are automatically used for server-side encryption by default.

NEW QUESTION 104
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
Scale down the application to one larger EC2 instance where only one instance is recording logs
6. 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 106
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 107
A developer wants to deploy a new version of an AWS Elastic Beanstalk application. During deployment, the application must maintain full capacity and avoid service interruption. Additionally, the developer must minimize the cost of additional resources that support the deployment.
Which deployment method should the developer use to meet these requirements?

A.

All at once
B. Rolling with additional batch
C. Blue/green
D. Immutable

Answer: D

Explanation:
The immutable deployment method is the best option for this scenario, because it meets the requirements of maintaining full capacity, avoiding service interruption, and minimizing the cost of additional resources.
The immutable deployment method creates a new set of instances in a separate Auto Scaling group and deploys the new version of the application to them. Then, it swaps the new instances with the old ones and terminates the old instances. This way, the application maintains full capacity during the deployment and avoids any downtime. The cost of additional resources is also minimized, because the new instances are only created for a short time and then replaced by the old ones.
The other deployment methods do not meet all the requirements:
? The all at once method deploys the new version to all instances simultaneously, which causes a short period of downtime and reduced capacity.
? The rolling with additional batch method deploys the new version in batches, but for the first batch it creates new instances instead of using the existing ones. This increases the cost of additional resources and reduces the capacity of the original environment.
? The blue/green method creates a new environment with a new set of instances and deploys the new version to them. Then, it swaps the URLs between the old and new environments. This method maintains full capacity and avoids service interruption, but it also increases the cost of additional resources significantly, because it duplicates the entire environment.

NEW QUESTION 108

A developer is creating a serverless application that uses an AWS Lambda function The developer will use AWS CloudFormation to deploy the application The application will write logs to Amazon CloudWatch Logs The developer has created a log group in a CloudFormation template for the application to use The developer needs to modify the CloudFormation template to make the name of the log group available to the application at runtime
Which solution will meet this requirement?

A. Use the AWS:Include transform in CloudFormation to provide the log group's name to the application
B. Pass the log group's name to the application in the user data section of the CloudFormation template.
C. Use the CloudFormation template's Mappings section to specify the log group's name for the application.
D. Pass the log group's Amazon Resource Name (ARN) as an environment variable to the Lambda function

Answer: D

Explanation:
FunctionName: MyLambdaFunction Code:
S3Bucket: your-lambda-code-bucket S3Key: lambda-code.zip
Runtime: nodejs14.x # Specify the desired runtime for your Lambda function Environment:
Variables:
LOG_GROUP_NAME: !Ref MyLogGroup https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-logs- loggroup.html

NEW QUESTION 109
A developer is investigating an issue in part of a company's application. In the application messages are sent to an Amazon Simple Queue Service (Amazon SQS) queue The AWS Lambda function polls messages from the SQS queue and sends email messages by using Amazon Simple Email Service (Amazon SES) Users have been receiving duplicate email messages during periods of high traffic.
Which reasons could explain the duplicate email messages? (Select TWO.)

A. Standard SQS queues support at-least-once message delivery
B. Standard SQS queues support exactly-once processing, so the duplicate email messages are because of user error.
C. Amazon SES has the DomainKeys Identified Mail (DKIM) authentication incorrectly configured
D. The SQS queue's visibility timeout is lower than or the same as the Lambda function's timeout.
E. The Amazon SES bounce rate metric is too high.

Answer: AD

Explanation:
Standard SQS queues support at-least-once message delivery, which means that a message can be delivered more than once to the same or different consumers. This can happen if the message is not deleted from the queue before the visibility timeout expires, or if there is a network issue or a system failure. The SQS queue's visibility timeout is the period of time that a message is invisible to other consumers after it is received by one consumer. If the visibility timeout is lower than or the same as the Lambda function's timeout, the Lambda function might not be able to process and delete the message before it becomes visible again, leading to duplicate processing and email messages. To avoid this, the visibility timeout should be set to at least 6 times the length of the Lambda function's timeout. The other options are not related to the issue of duplicate email messages. References
? Using the Amazon SQS message deduplication ID
? Exactly-once processing - Amazon Simple Queue Service
? Amazon SQS duplicated messages in queue - Stack Overflow
? amazon web services - How long can duplicate SQS messages persist …
? Standard SQS - Duplicate message | AWS re:Post - Amazon Web Services, Inc.

NEW QUESTION 114
A developer is building a microservices-based application by using Python on AWS and several AWS services The developer must use AWS X-Ray The developer views the service map by using the console to view the service dependencies. During testing, the developer notices that some services are missing from the service map
What can the developer do to ensure that all services appear in the X-Ray service map?

A. Modify the X-Ray Python agent configuration in each service to increase the sampling rate
B. Instrument the application by using the X-Ray SDK for Pytho
C. Install the X-Ray SDK for all the services that the application uses
D. Enable X-Ray data aggregation in Amazon CloudWatch Logs for all the services that the application uses
E. Increase the X-Ray service map timeout value in the X-Ray console

Answer: B

Explanation:
The X-Ray SDK for Python provides libraries and tools for instrumenting Python applications that use AWS services and other AWS X-Ray integrations. By installing the X-Ray SDK for all the services that the application uses, the developer can ensure that all the service dependencies are captured and displayed in the X-Ray service map. The other options are not relevant or effective for this scenario. References
? AWS X-Ray SDK for Python
? Instrumenting a Python Application

NEW QUESTION 115
A company has deployed an application on AWS Elastic Beanstalk. The company has configured the Auto Scaling group that is associated with the Elastic Beanstalk environment to have five Amazon EC2 instances. If the capacity is fewer than four EC2 instances during the deployment, application performance degrades. The company is using the all-at-once deployment policy.
What is the MOST cost-effective way to solve the deployment issue?

A. Change the Auto Scaling group to six desired instances.
B. Change the deployment policy to traffic splittin
C. Specify an evaluation time of 1 hour.
D. Change the deployment policy to rolling with additional batc
E. Specify a batch size of 1.
F. Change the deployment policy to rollin
G. Specify a batch size of 2.

Answer: C

Explanation:
This solution will solve the deployment issue by deploying the new version of the application to one new EC2 instance at a time, while keeping the old version running on

the existing instances. This way, there will always be at least four instances serving traffic during the deployment, and no downtime or performance degradation will occur. Option A is not optimal because it will increase the cost of running the Elastic Beanstalk environment without solving the deployment issue. Option B is not optimal because it will split the traffic between two versions of the application, which may cause inconsistency and confusion for the customers. Option D is not optimal because it will deploy the new version of the application to two existing instances at a time, which may reduce the capacity below four instances during the deployment.
References: AWS Elastic Beanstalk Deployment Policies

NEW QUESTION 116

A company has an Amazon S3 bucket that contains sensitive data. The data must be encrypted in transit and at rest. The company encrypts the data in the S3 bucket by using an AWS Key Management Service (AWS KMS) key. A developer needs to grant several other AWS accounts the permission to use the S3 GetObject operation to retrieve the data from the S3 bucket.
How can the developer enforce that all requests to retrieve the data provide encryption in transit?

A. Define a resource-based policy on the S3 bucket to deny access when a request meets the condition “aws:SecureTransport”: “false”.
B. Define a resource-based policy on the S3 bucket to allow access when a request meets the condition “aws:SecureTransport”: “false”.
C. Define a role-based policy on the other accounts' roles to deny access when a request meets the condition of “aws:SecureTransport”: “false”.
D. Define a resource-based policy on the KMS key to deny access when a request meets the condition of “aws:SecureTransport”: “false”.

Answer: A

Explanation:
Amazon S3 supports resource-based policies, which are JSON documents that specify the permissions for accessing S3 resources. A resource-based policy can be used to enforce encryption in transit by denying access to requests that do not use HTTPS. The condition key aws:SecureTransport can be used to check if the request was sent using SSL. If the value of this key is false, the request is denied; otherwise, the request is allowed. Reference: How do I use an S3 bucket policy to require requests to use Secure Socket Layer (SSL)?

NEW QUESTION 119
A company is preparing to migrate an application to the company's first AWS environment Before this migration, a developer is creating a proof-of-concept application to validate a model for building and deploying container-based applications on AWS.
Which combination of steps should the developer take to deploy the containerized proof-of- concept application with the LEAST operational effort? (Select TWO.)

A. Mastered
B. Not Mastered

Answer: A

Explanation:
To deploy a containerized application on AWS with the least operational effort, the developer should package the application into a container image by using the Docker CLI and upload the image to Amazon ECR, which is a fully managed container registry service. Then, the developer should deploy the application to Amazon ECS on AWS Fargate, which is a serverless compute engine for containers that eliminates the need to provision and manage servers or clusters. Amazon ECS will automatically scale, load balance, and monitor the application. References
? How to Deploy Docker Containers | AWS
? Deploy a Web App Using AWS App Runner
? How to Deploy Containerized Apps on AWS Using ECR and Docker

NEW QUESTION 121
A developer is building a serverless application that is based on AWS Lambda. The developer initializes the AWS software development kit (SDK) outside of the Lambda handcar function.
What is the PRIMARY benefit of this action?

A. Improves legibility and systolic convention
B. Takes advantage of runtime environment reuse
C. Provides better error handling
D. Creates a new SDK instance for each invocation

Answer: B

Explanation:
This benefit occurs when initializing the AWS SDK outside of the Lambda handler function because it allows the SDK instance to be reused across multiple invocations of the same function. This can improve performance and reduce latency by avoiding unnecessary initialization overhead. If the SDK is initialized inside the handler function, it will create a new SDK instance for each invocation, which can increase memory usage and execution time.
Reference: [AWS Lambda execution environment], [Best Practices for Working with AWS
Lambda Functions]

NEW QUESTION 122
A company has a social media application that receives large amounts of traffic User posts and interactions are continuously updated in an Amazon RDS database The data changes frequently, and the data types can be complex The application must serve read requests with minimal latency
The application's current architecture struggles to deliver these rapid data updates efficiently The company needs a solution to improve the application's performance.
Which solution will meet these requirements'?

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Creating an Amazon ElastiCache for Redis cluster is the best solution for improving the application’s performance. Redis is an in-memory data store that can serve read requests with minimal latency and handle complex data types, such as lists, sets, hashes, and streams. By using a write-through caching strategy, the application can ensure that the data in Redis is always consistent with the data in RDS. The application can read the data from Redis instead of RDS, reducing the load on the database and improving the response time. The other solutions are either not feasible or not effective. Amazon DynamoDB Accelerator (DAX) is a caching service that works only with DynamoDB, not RDS. Amazon S3 Transfer Acceleration is a feature that speeds up data transfers between S3 and clients across the internet, not between RDS and the application. Amazon CloudFront is a content delivery network that can cache static content, such as images, videos, or HTML files, but not dynamic content, such as user posts and
interactions. References
? Amazon ElastiCache for Redis
? Caching Strategies and Best Practices - Amazon ElastiCache for Redis
? Using Amazon ElastiCache for Redis with Amazon RDS
? Amazon DynamoDB Accelerator (DAX)
? Amazon S3 Transfer Acceleration
? Amazon CloudFront

NEW QUESTION 125

When a developer tries to run an AWS Code Build project, it raises an error because the length of all environment variables exceeds the limit for the combined maximum of characters.
What is the recommended solution?

A. Add the export LC-_ALL" on _ US, tuft" command to the pre _ build section to ensure POSIX Localization.
B. Use Amazon Cognate to store key-value pairs for large numbers of environment variables
C. Update the settings for the build project to use an Amazon S3 bucket for large numbers of environment variables
D. Use AWS Systems Manager Parameter Store to store large numbers ot environment variables

Answer: D

Explanation:
This solution allows the developer to overcome the limit for the combined maximum of characters for environment variables in AWS CodeBuild. AWS Systems Manager Parameter Store provides secure, hierarchical storage for configuration data management and secrets management. The developer can store large numbers of environment variables as parameters in Parameter Store and reference them in the buildspec file using parameter references. Adding export LC_ALL="en_US.utf8" command to the pre_build section will not affect the environment variables limit. Using Amazon Cognito or an Amazon S3 bucket to store key-value pairs for environment variables will require additional configuration and integration.
Reference: [Build Specification Reference for AWS CodeBuild], [What Is AWS Systems
Manager Parameter Store?]

NEW QUESTION 127
A company needs to distribute firmware updates to its customers around the world.
Which service will allow easy and secure control of the access to the downloads at the lowest cost?

A. Use Amazon CloudFront with signed URLs for Amazon S3.
B. Create a dedicated Amazon CloudFront Distribution for each customer.
C. Use Amazon CloudFront with AWS Lambda@Edge.
D. Use Amazon API Gateway and AWS Lambda to control access to an S3 bucket.

Answer: A

Explanation:
This solution allows easy and secure control of access to the downloads at the lowest cost because it uses a content delivery network (CDN) that can cache and distribute firmware updates to customers around the world, and uses a mechanism that can restrict access to specific files or versions. Amazon CloudFront is a CDN that can improve performance, availability, and security of web applications by delivering content from edge locations closer to customers. Amazon S3 is a storage service that can store firmware updates in buckets and objects. Signed URLs are URLs that include additional information, such as an expiration date and time, that give users temporary access to specific objects in S3 buckets. The developer can use CloudFront to serve firmware updates from S3 buckets and use signed URLs to control who can download them and for how long. Creating a dedicated CloudFront distribution for each customer will incur unnecessary costs and complexity. Using Amazon CloudFront with AWS Lambda@Edge will require additional programming overhead to implement custom logic at the edge locations. Using Amazon API Gateway and AWS Lambda to control access to an S3 bucket will also require additional programming overhead and may not provide optimal performance or availability.
Reference: [Serving Private Content through CloudFront], [Using CloudFront with Amazon
S3]

NEW QUESTION 128
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 129
An AWS Lambda function requires read access to an Amazon S3 bucket and requires read/write access to an Amazon DynamoDB table The correct 1AM policy already exists
What is the MOST secure way to grant the Lambda function access to the S3 bucket and the DynamoDB table?

A. Attach the existing 1AM policy to the Lambda function.
B. Create an 1AM role for the Lambda function Attach the existing 1AM policy to the role Attach the role to the Lambda function
C. Create an 1AM user with programmatic access Attach the existing 1AM policy to the use
D. Add the user access key ID and secret access key as environment variables in the Lambda function.
E. Add the AWS account root user access key ID and secret access key as encrypted environment variables in the Lambda function

Answer: B

Explanation:
The most secure way to grant the Lambda function access to the S3 bucket and the DynamoDB table is to create an IAM role for the Lambda function and attach the existing IAM policy to the role. This way, you can use the principle of least privilege and avoid exposing any credentials in your function code or environment variables. You can also leverage the temporary security credentials that AWS provides to the Lambda function when it assumes the role. This solution follows the best practices for working with AWS Lambda functions1 and designing and architecting with DynamoDB2. References
? Best practices for working with AWS Lambda functions
? Best practices for designing and architecting with DynamoDB

NEW QUESTION 130
An application runs on multiple EC2 instances behind an ELB.
Where is the session data best written so that it can be served reliably across multiple requests?

A. Write data to Amazon ElastiCache
B. Write data to Amazon Elastic Block Store
C. Write data to Amazon EC2 instance Store
D. Wide data to the root filesystem

Answer: A

Explanation:
The solution that will meet the requirements is to write data to Amazon ElastiCache. This way, the application can write session data to a fast, scalable, and reliable in-memory data store that can be served reliably across multiple requests. The other options either involve writing data to persistent storage, which is slower and more expensive than in-memory storage, or writing data to the root filesystem, which is not shared among multiple EC2 instances.
Reference: Using ElastiCache for session management

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 134
A developer must analyze performance issues with production-distributed applications written as AWS Lambda functions. These distributed Lambda applications invoke other components that make up me applications. How should the developer identify and troubleshoot the root cause of the performance issues in production?

A. Add logging statements to the Lambda function
B. then use Amazon CloudWatch to view the logs.
C. Use AWS CloudTrail and then examine the logs.
D. Use AWS X-Ra
E. then examine the segments and errors.
F. Run Amazon inspector agents and then analyze performance.

Answer: C

Explanation:
This solution will meet the requirements by using AWS X-Ray to analyze and debug the performance issues with the distributed Lambda applications. AWS X-Ray is a service that collects data about requests that the applications serve, and provides tools to view, filter, and gain insights into that data. The developer can use AWS X-Ray to identify the root cause of the performance issues by examining the segments and errors that show the details of each request and the components that make up the applications. Option A is not optimal because it will use logging statements and Amazon CloudWatch, which may not provide enough information or visibility into the distributed applications. Option B is not optimal because it will use AWS CloudTrail, which is a service that records API calls and events for AWS services, not application performance data. Option D is not optimal because it will use Amazon Inspector, which is a service that helps improve the security and compliance of applications on Amazon EC2 instances, not Lambda functions. References: AWS X-Ray, Using AWS X-Ray with AWS Lambda

NEW QUESTION 138
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 StringNotEquals condition operator is used to negate the condition, so that only requests from the listed VPC endpoints are allowed.
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 143
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 148
An developer is building a serverless application by using the AWS Serverless Application Model (AWS SAM). The developer is currently testing the application in a development environment. When the application is nearly finsihed, the developer will need to set up additional testing and staging environments for a quality assurance team.
The developer wants to use a feature of the AWS SAM to set up deployments to multiple environments.
Which solution will meet these requirements with the LEAST development effort?

A. Add a configuration file in TOML format to group configuration entries to every environmen
B. Add a table for each testing and staging environmen
C. Deploy updates to the environments by using the sam deploy command and the --config-env flag that corresponds to the each environment.
D. Create additional AWS SAM templates for each testing and staging environmen
E. Write a custom shell script that uses the sam deploy command and the --template-file flag to deploy updates to the environments.
F. Create one AWS SAM configuration file that has default parameter
G. Perform updates to the testing and staging environments by using the —parameter-overrides flag in the AWS SAM CLI and the parameters that the updates will override.
H. Use the existing AWS SAM templat
I. Add additional parameters to configure specific attributes for the serverless function and database table resources that are in each environmen
J. Deploy updates to the testing and staging environments by using the sam deploy command.

Answer: A

Explanation:
The correct answer is A. Add a configuration file in TOML format to group configuration entries to every environment. Add a table for each testing and staging environment. Deploy updates to the environments by using the sam deploy command and the --config-env flag that corresponds to the each environment.
* A. Add a configuration file in TOML format to group configuration entries to every environment. Add a table for each testing and staging environment. Deploy updates to the environments by using the sam deploy command and the --config-env flag that corresponds to the each environment. This is correct. This solution will meet the requirements with the least development effort, because it uses a feature of the AWS SAM CLI that supports a project-level configuration file that can be used to configure AWS SAM CLI command parameter values1. The configuration file can have multiple environments, each with its own set of parameter values, such as stack name, region, capabilities, and more2. The developer can use the --config-env option to specify which environment to use when deploying the application3. This way, the developer can avoid creating multiple templates or scripts, or manually overriding parameters for each environment.
* B. Create additional AWS SAM templates for each testing and staging environment. Write a custom shell script that uses the sam deploy command and the --template-file flag to deploy updates to the environments. This is incorrect. This solution will not meet the requirements with the least development effort, because it requires creating and maintaining multiple templates and scripts for each environment. This can introduce duplication, inconsistency, and complexity in the deployment process.
* C. Create one AWS SAM configuration file that has default parameters. Perform updates to the testing and staging environments by using the —parameter-overrides flag in the AWS SAM CLI and the parameters that the updates will override. This is incorrect. This solution will not meet the requirements with the least development effort, because it requires manually specifying and overriding parameters for each environment every time the developer deploys the application. This can be error-prone, tedious, and inefficient.
* D. Use the existing AWS SAM template. Add additional parameters to configure specific attributes for the serverless function and database table resources that are in each environment. Deploy updates to the testing and staging environments by using the sam deploy command. This is incorrect. This solution will not meet the requirements with the least development effort, because it requires modifying the existing template and adding complexity to the resource definitions for each environment. This can also make it difficult to manage and track changes across different environments.
References:
? 1: AWS SAM CLI configuration file - AWS Serverless Application Model
? 2: Configuration file basics - AWS Serverless Application Model
? 3: Specify a configuration file - AWS Serverless Application Model

NEW QUESTION 149
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 152
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 157
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 160
A development team maintains a web application by using a single AWS CloudFormation template. The template defines web servers and an Amazon RDS database. The team uses the Cloud Formation template to deploy the Cloud Formation stack to different environments.
During a recent application deployment, a developer caused the primary development database to be dropped and recreated. The result of this incident was a loss of data. The team needs to avoid accidental database deletion in the future.
Which solutions will meet these requirements? (Choose two.)

A. Add a CloudFormation Deletion Policy attribute with the Retain value to the database resource.
B. Update the CloudFormation stack policy to prevent updates to the database.
C. Modify the database to use a Multi-AZ deployment.
D. Create a CloudFormation stack set for the web application and database deployments.
E. Add a Cloud Formation DeletionPolicy attribute with the Retain value to the stack.

Answer: AB

Explanation:
AWS CloudFormation is a service that enables developers to model and provision AWS resources using templates. The developer can add a CloudFormation Deletion Policy attribute with the Retain value to the database resource. This will prevent the database from being deleted when the stack is deleted or updated. The developer can also update the CloudFormation stack policy to prevent updates to the database. This will prevent accidental changes to the database configuration or properties.

References:
? [What Is AWS CloudFormation? - AWS CloudFormation]
? [DeletionPolicy Attribute - AWS CloudFormation]
? [Protecting Resources During Stack Updates - AWS CloudFormation]

NEW QUESTION 165
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

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