Amazon

Exam Questions DVA-C02

DVA-C02

NEW QUESTION 1
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 2
A developer is building an application that uses AWS API Gateway APIs. AWS Lambda function, and AWS Dynamic DB tables. The developer uses the AWS Serverless Application Model (AWS SAM) to build and run serverless applications on AWS. Each time the developer pushes of changes for only to the Lambda functions, all the artifacts in the application are rebuilt.
The developer wants to implement AWS SAM Accelerate by running a command to only redeploy the Lambda functions that have changed.
Which command will meet these requirements?

A. sam deploy -force-upload
B. sam deploy -no-execute-changeset
C. sam package
D. sam sync -watch

Answer: D

Explanation:
The command that will meet the requirements is sam sync -watch. This command enables AWS SAM Accelerate mode, which allows the developer to only redeploy the Lambda functions that have changed. The -watch flag enables file watching, which automatically detects changes in the source code and triggers a redeployment. The other commands either do not enable AWS SAM Accelerate mode, or do not redeploy the Lambda functions automatically.
Reference: AWS SAM Accelerate

NEW QUESTION 3
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 4
A developer is designing a serverless application for a game in which users register and log in through a web browser The application makes requests on behalf of users to a set of AWS Lambda functions that run behind an Amazon API Gateway HTTP API
The developer needs to implement a solution to register and log in users on the application's sign-in page. The solution must minimize operational overhead and must minimize ongoing management of user identities.
Which solution will meet these requirements'?

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

Answer: A

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

NEW QUESTION 5
A company is building a serverless application on AWS. The application uses an AWS Lambda function to process customer orders 24 hours a day, 7 days a week. The Lambda function calls an external vendor's HTTP API to process payments.
During load tests, a developer discovers that the external vendor payment processing API occasionally times out and returns errors. The company expects that some payment processing API calls will return errors.
The company wants the support team to receive notifications in near real time only when the payment processing external API error rate exceed 5% of the total number of transactions in an hour. Developers need to use an existing Amazon Simple Notification Service (Amazon SNS) topic that is configured to notify the support team.
Which solution will meet these requirements?

A. Write the results of payment processing API calls to Amazon CloudWatc
B. Use Amazon CloudWatch Logs Insights to query the CloudWatch log
C. Schedule the Lambda function to check the CloudWatch logs and notify the existing SNS topic.
D. Publish custom metrics to CloudWatch that record the failures of the external payment processing API call
E. Configure a CloudWatch alarm to notify the existing SNS topic when error rate exceeds the specified rate.
F. Publish the results of the external payment processing API calls to a new Amazon SNS topi
G. Subscribe the support team members to the new SNS topic.
H. Write the results of the external payment processing API calls to Amazon S3. Schedule an Amazon Athena query to run at regular interval
I. Configure Athena to send notifications to the existing SNS topic when the error rate exceeds the specified rate.

Answer: B

Explanation:
Amazon CloudWatch is a service that monitors AWS resources and applications. The developer can publish custom metrics to CloudWatch that record the failures of the external payment processing API calls. The developer can configure a CloudWatch alarm to notify the existing SNS topic when the error rate exceeds 5% of the total number of transactions in an hour. This solution will meet the requirements in a near real-time and scalable way.
References:
? [What Is Amazon CloudWatch? - Amazon CloudWatch]
? [Publishing Custom Metrics - Amazon CloudWatch]
? [Creating Amazon CloudWatch Alarms - Amazon CloudWatch]

NEW QUESTION 6
A company is building an application for stock trading. The application needs sub- millisecond latency for processing trade requests. The company uses Amazon DynamoDB to store all the trading data that is used to process each trading request A development team performs load testing on the application and finds that the data retrieval time is higher than expected. The development team needs a solution that reduces the data retrieval time with the least possible effort.
Which solution meets these requirements'?

A. Add local secondary indexes (LSis) for the trading data.
B. Store the trading data m Amazon S3 and use S3 Transfer Acceleration.
C. Add retries with exponential back off for DynamoDB queries.
D. Use DynamoDB Accelerator (DAX) to cache the trading data.

Answer: D

Explanation:
This solution will meet the requirements by using DynamoDB Accelerator (DAX), which is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10 times performance improvement - from milliseconds to microseconds - even at millions of requests per second. The developer can use DAX to cache the trading data that is used to process each trading request, which will reduce the data retrieval time with the least possible effort. Option A is not optimal because it will add local secondary indexes (LSIs) for the trading data, which may not improve the performance or reduce the latency of data retrieval, as LSIs are stored on the same partition as the base table and share the same provisioned throughput. Option B is not optimal because it will store the trading data in Amazon S3 and use S3 Transfer Acceleration, which is a feature that enables fast, easy, and secure transfers of files over long distances between S3 buckets and clients, not between DynamoDB and clients. Option C is not optimal because it will add retries with exponential backoff for DynamoDB queries, which is a strategy to handle transient errors by retrying failed requests with increasing delays, not by reducing data retrieval time.
References: [DynamoDB Accelerator (DAX)], [Local Secondary Indexes]

NEW QUESTION 7
A developer needs to migrate an online retail application to AWS to handle an anticipated increase in traffic. The application currently runs on two servers: one server for the web application and another server for the database. The web server renders webpages and manages session state in memory. The database server hosts a MySQL database that contains order details. When traffic to the application is heavy, the memory usage for the web server approaches 100% and the application slows down considerably.
The developer has found that most of the memory increase and performance decrease is related to the load of managing additional user sessions. For the web server migration, the developer will use Amazon EC2 instances with an Auto Scaling group behind an Application Load Balancer.
Which additional set of changes should the developer make to the application to improve the application's performance?

A. Use an EC2 instance to host the MySQL databas
B. Store the session data and the application data in the MySQL database.
C. Use Amazon ElastiCache for Memcached to store and manage the session dat
D. Use an Amazon RDS for MySQL DB instance to store the application data.
E. Use Amazon ElastiCache for Memcached to store and manage the session data and the application data.
F. Use the EC2 instance store to manage the session dat
G. Use an Amazon RDS for MySQL DB instance to store the application data.

Answer: B

Explanation:
Using Amazon ElastiCache for Memcached to store and manage the session data will reduce the memory load and improve the performance of the web server. Using Amazon RDS for MySQL DB instance to store the application data will provide a scalable, reliable, and managed database service. Option A is not optimal because it does not address the memory issue of the web server. Option C is not optimal because it does not provide a persistent storage for the application data. Option D is not optimal because it does not provide a high availability and durability for the session data.
References: Amazon ElastiCache, Amazon RDS

NEW QUESTION 8
A developer is testing a new file storage application that uses an Amazon CloudFront distribution to serve content from an Amazon S3 bucket. The distribution accesses the S3 bucket by using an origin access identity (OAI). The S3 bucket's permissions explicitly deny access to all other users. The application prompts users to authenticate on a login page and then uses signed cookies to allow users to access their personal storage directories. The developer has configured the distribution to use its default cache behavior with restricted viewer access and has set the origin to point to the S3 bucket. However, when the developer tries to navigate to the login page, the developer receives a 403 Forbidden error.
The developer needs to implement a solution to allow unauthenticated access to the login page. The solution also must keep all private content secure.
Which solution will meet these requirements?

A. Add a second cache behavior to the distribution with the same origin as the default cache behavio
B. Set the path pattern for the second cache behavior to the path of the login page, and make viewer access unrestricte
C. Keep the default cache behavior's settings unchanged.
D. Add a second cache behavior to the distribution with the same origin as the default cache behavio
E. Set the path pattern for the second cache behavior to *, and make viewer access restricte
F. Change the default cache behavior's path pattern to the path of the login page, and make viewer access unrestricted.
G. Add a second origin as a failover origin to the default cache behavio
H. Point the failover origin to the S3 bucke
I. Set the path pattern for the primary origin to *, and make viewer access restricte
J. Set the path pattern for the failover origin to the path of the login page, and make viewer access unrestricted.
K. Add a bucket policy to the S3 bucket to allow read acces
L. Set the resource on the policy to the Amazon Resource Name (ARN) of the login page object in the S3 bucke
M. Add a CloudFront function to the default cache behavior to redirect unauthorized requests to the login page's S3 URL.

Answer: A

Explanation:
The solution that will meet the requirements is to add a second cache behavior to the distribution with the same origin as the default cache behavior. Set the path pattern for the second cache behavior to the path of the login page, and make viewer access unrestricted. Keep the default cache behavior’s settings unchanged. This way, the login page can be accessed without authentication, while all other content remains secure and requires signed cookies. The other options either do not allow unauthenticated access to the login page, or expose private content to unauthorized users.
Reference: Restricting Access to Amazon S3 Content by Using an Origin Access Identity

NEW QUESTION 9
A company needs to deploy all its cloud resources by using AWS CloudFormation templates A developer must create an Amazon Simple Notification Service (Amazon SNS) automatic notification to help enforce this rule. The developer creates an SNS topic and subscribes the email address of the company's security team to the SNS topic.
The security team must receive a notification immediately if an 1AM role is created without the use of CloudFormation.
Which solution will meet this requirement?

A. Create an AWS Lambda function to filter events from CloudTrail if a role was created without CloudFormation Configure the Lambda function to publish to the SNS topi
B. Create an Amazon EventBridge schedule to invoke the Lambda function every 15 minutes
C. Create an AWS Fargate task in Amazon Elastic Container Service (Amazon ECS) to filter events from CloudTrail if a role was created without CloudFormation Configure the Fargate task to publish to the SNS topic Create an Amazon EventBridge schedule to run the Fargate task every 15 minutes
D. Launch an Amazon EC2 instance that includes a script to filter events from CloudTrail if a role was created without CloudFormatio
E. Configure the script to publish to the SNS topi
F. Create a cron job to run the script on the EC2 instance every 15 minutes.
G. Create an Amazon EventBridge rule to filter events from CloudTrail if a role was created without CloudFormation Specify the SNS topic as the target of the EventBridge rule.

Answer: D

Explanation:
Creating an Amazon EventBridge rule is the most efficient and scalable way to monitor and react to events from CloudTrail, such as the creation of an IAM role without CloudFormation. EventBridge allows you to specify a filter pattern to match the events you are interested in, and then specify an SNS topic as the target to send notifications. This solution does not require any additional resources or code, and it can trigger notifications in near real-time. The other solutions involve creating and managing additional resources, such as Lambda functions, Fargate tasks, or EC2 instances, and they rely on polling CloudTrail events every 15 minutes, which can introduce delays and increase
costs. References
? Using Amazon EventBridge rules to process AWS CloudTrail events
? Using AWS CloudFormation to create and manage AWS Batch resources
? How to use AWS CloudFormation to configure auto scaling for Amazon Cognito and AWS AppSync
? Using AWS CloudFormation to automate the creation of AWS WAF web ACLs, rules, and conditions

NEW QUESTION 10
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 10
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 11
A company has an application that is hosted on Amazon EC2 instances The application stores objects in an Amazon S3 bucket and allows users to download objects from the S3 bucket A developer turns on S3 Block Public Access for the S3 bucket After this change, users report errors when they attempt to download objects The developer needs to implement a solution so that only users who are signed in to the application can access objects in the S3 bucket.
Which combination of steps will meet these requirements in the MOST secure way? (Select TWO.)

A. Create an EC2 instance profile and role with an appropriate policy Associate the role with the EC2 instances
B. Create an 1AM user with an appropriate polic
C. Store the access key ID and secret access key on the EC2 instances
D. Modify the application to use the S3 GeneratePresignedUrl API call
E. Modify the application to use the S3 GetObject API call and to return the object handle to the user
F. Modify the application to delegate requests to the S3 bucket.

Answer: AC

Explanation:
The most secure way to allow the EC2 instances to access the S3 bucket is to use an EC2 instance profile and role with an appropriate policy that grants the necessary permissions. This way, the EC2 instances can use temporary security credentials that are automatically rotated and do not need to store any access keys on the instances. To allow the users who are signed in to the application to download objects from the S3 bucket, the application can use the S3 GeneratePresignedUrl API call to create a pre-signed URL that grants temporary access to a specific object. The pre-signed URL can be returned to the user, who can then use it to download the object within a specified time period. References
? Use Amazon S3 with Amazon EC2
? How to Access AWS S3 Bucket from EC2 Instance In a Secured Way
? Sharing an Object with Others

NEW QUESTION 12
A company has a web application that is hosted on Amazon EC2 instances The EC2 instances are configured to stream logs to Amazon CloudWatch Logs The company needs to receive an Amazon Simple Notification Service (Amazon SNS) notification when the number of application error messages exceeds a defined threshold within a 5-minute period
Which solution will meet these requirements?

A. Rewrite the application code to stream application logs to Amazon SNS Configure an SNS topic to send a notification when the number of errors exceeds the defined threshold within a 5-minute period
B. Configure a subscription filter on the CloudWatch Logs log grou
C. Configure the filter to send an SNS notification when the number of errors exceeds the defined threshold within a 5-minute period.
D. Install and configure the Amazon Inspector agent on the EC2 instances to monitor for errors Configure Amazon Inspector to send an SNS notification when the number of errors exceeds the defined threshold within a 5-minute period
E. Create a CloudWatch metric filter to match the application error pattern in the log data. Set up a CloudWatch alarm based on the new custom metri
F. Configure the alarm to send an SNS notification when the number of errors exceeds the defined threshold within a 5- minute period.

Answer: D

Explanation:
The best solution is to create a CloudWatch metric filter to match the application error pattern in the log data. This will allow you to create a custom metric that tracks the number of errors in your application. You can then set up a CloudWatch alarm based on this metric and configure it to send an SNS notification when the number of errors exceeds a defined threshold within a 5-minute period. This solution does not require any changes to your application code or installing any additional agents on your EC2 instances. It also leverages the existing integration between CloudWatch and SNS for sending notifications. References
? Create Metric Filters - Amazon CloudWatch Logs
? Creating Amazon CloudWatch Alarms - Amazon CloudWatch
? How to send alert based on log message on CloudWatch - Stack Overflow

NEW QUESTION 17
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 18
A company is using an AWS Lambda function to process records from an Amazon Kinesis data stream. The company recently observed slow processing of the records. A developer notices that the iterator age metric for the function is increasing and that the Lambda run duration is constantly above normal.
Which actions should the developer take to increase the processing speed? (Choose two.)

A. Increase the number of shards of the Kinesis data stream.
B. Decrease the timeout of the Lambda function.
C. Increase the memory that is allocated to the Lambda function.
D. Decrease the number of shards of the Kinesis data stream.
E. Increase the timeout of the Lambda function.

Answer: AC

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

NEW QUESTION 21
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 22
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.
Store credentials in a config fil
J.
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 26
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 31
A company is building a web application on AWS. When a customer sends a request, the application will generate reports and then make the reports available to the customer within one hour. Reports should be accessible to the customer for 8 hours. Some reports are larger than 1 MB. Each report is unique to the customer. The application should delete all reports that are older than 2 days.
Which solution will meet these requirements with the LEAST operational overhead?

A. Generate the reports and then store the reports as Amazon DynamoDB items that have a specified TT
B. Generate a URL that retrieves the reports from DynamoD
C. Provide the URL to customers through the web application.
D. Generate the reports and then store the reports in an Amazon S3 bucket that uses server-side encryptio
E. Attach the reports to an Amazon Simple Notification Service (Amazon SNS) messag
F. Subscribe the customer to email notifications from Amazon SNS.
G. Generate the reports and then store the reports in an Amazon S3 bucket that uses server-side encryptio
H. Generate a presigned URL that contains an expiration date Provide the URL to customers through the web applicatio
I. Add S3 Lifecycle configuration rules to the S3 bucket to delete old reports.
J. Generate the reports and then store the reports in an Amazon RDS database with a date stam
K. Generate an URL that retrieves the reports from the RDS databas
L. Provide the URL to customers through the web applicatio
M. Schedule an hourly AWS Lambda function to delete database records that have expired date stamps.

Answer: C

Explanation:
This solution will meet the requirements with the least operational overhead because it uses Amazon S3 as a scalable, secure, and durable storage service for the reports. The presigned URL will allow customers to access their reports for a limited time (8 hours) without requiring additional authentication. The S3 Lifecycle configuration rules will automatically delete the reports that are older than 2 days, reducing storage costs and complying with the data retention policy. Option A is not optimal because it will incur additional costs and complexity to store the reports as DynamoDB items, which have a size limit of 400 KB. Option B is not optimal because it will not provide customers with access to their reports within one hour, as Amazon SNS email delivery is not guaranteed. Option D is not optimal because it will require more operational overhead to manage an RDS database and a Lambda function for storing and deleting the reports.
References: Amazon S3 Presigned URLs, Amazon S3 Lifecycle

NEW QUESTION 32
A developer is building a serverless application by using AWS Serverless Application Model (AWS SAM) on multiple AWS Lambda functions. When the application is deployed, the developer wants to shift 10% of the traffic to the new deployment of the application for the first 10 minutes after deployment. If there are no issues, all traffic must switch over to the new version.
Which change to the AWS SAM template will meet these requirements?

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

Answer: A

Explanation:
? The Deployment Preference Type property specifies how traffic should be shifted between versions of a Lambda function1. The Canary10Percent10Minutes option means that 10% of the traffic is immediately shifted to the new version, and after 10 minutes, the remaining 90% of the traffic is shifted1. This matches the requirement of shifting 10% of the traffic for the first 10 minutes, and then switching all traffic to the new version.
? The AutoPublishAlias property enables AWS SAM to automatically create and update a Lambda alias that points to the latest version of the function1. This is required to use the Deployment Preference Type property1. The alias name can be specified by the developer, and it can be used to invoke the function with the latest code.

NEW QUESTION 35
A developer is working on a serverless application that needs to process any changes to an Amazon DynamoDB table with an AWS Lambda function.
How should the developer configure the Lambda function to detect changes to the DynamoDB table?

A. Create an Amazon Kinesis data stream, and attach it to the DynamoDB tabl
B. Create a trigger to connect the data stream to the Lambda function.
C. Create an Amazon EventBridge rule to invoke the Lambda function on a regular schedul
D. Connect to the DynamoDB table from the Lambda function to detect changes.
E. Enable DynamoDB Streams on the tabl

F. Create a trigger to connect the DynamoDB stream to the Lambda function.
G. Create an Amazon Kinesis Data Firehose delivery stream, and attach it to the DynamoDB tabl
H. Configure the delivery stream destination as the Lambda function.

Answer: C

Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance with seamless scalability. DynamoDB Streams is a feature that captures data modification events in DynamoDB tables. The developer can enable DynamoDB Streams on the table and create a trigger to connect the DynamoDB stream to the Lambda function. This solution will enable the Lambda function to detect changes to the DynamoDB table in near real time.
References:
? [Amazon DynamoDB]
? [DynamoDB Streams - Amazon DynamoDB]
? [Using AWS Lambda with Amazon DynamoDB - AWS Lambda]

NEW QUESTION 38
A developer has created an AWS Lambda function that makes queries to an Amazon Aurora MySQL DB instance. When the developer performs a test the OB instance shows an error for too many connections.
Which solution will meet these requirements with the LEAST operational effort?

A. Create a read replica for the DB instance Query the replica DB instance instead of the primary DB instance.
B. Migrate the data lo an Amazon DynamoDB database.
C. Configure the Amazon Aurora MySQL DB instance tor Multi-AZ deployment.
D. Create a proxy in Amazon RDS Proxy Query the proxy instead of the DB instance.

Answer: D

Explanation:
This solution will meet the requirements by using Amazon RDS Proxy, which is a fully managed, highly available database proxy for Amazon RDS that makes applications more scalable, more resilient to database failures, and more secure. The developer can create a proxy in Amazon RDS Proxy, which sits between the application
and the DB instance and handles connection management, pooling, and routing. The developer can query the proxy instead of the DB instance, which reduces the number of open connections to the DB instance and avoids errors for too many connections. Option A is not optimal because it will create a read replica for the DB instance, which may not solve the problem of too many connections as read replicas also have connection limits and may incur additional costs. Option B is not optimal because it will migrate the data to an Amazon DynamoDB database, which may introduce additional complexity and overhead for migrating and accessing data from a different database service. Option C is not optimal because it will configure the Amazon Aurora MySQL DB instance for Multi-AZ deployment, which may improve availability and durability of the DB instance but not reduce the number of connections.
References: [Amazon RDS Proxy], [Working with Amazon RDS Proxy]

NEW QUESTION 40
A company has an analytics application that uses an AWS Lambda function to process transaction data asynchronously A developer notices that asynchronous invocations of the Lambda function sometimes fail When failed Lambda function invocations occur, the developer wants to invoke a second Lambda function to handle errors and log details.
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Configuring a Lambda function destination with a failure condition is the best solution for invoking a second Lambda function to handle errors and log details. A Lambda function destination is a resource that Lambda sends events to after a function is invoked. The developer can specify the destination type as Lambda function and the ARN of the error-handling Lambda function as the resource. The developer can also specify the failure condition, which means that the destination is invoked only when the initial Lambda function fails. The destination event will include the response from the initial function, the request ID, and the timestamp. The other solutions are either not feasible or not efficient. Enabling AWS X-Ray active tracing on the initial Lambda function will help to monitor and troubleshoot the function performance, but it will not automatically invoke the error-handling Lambda function. Configuring a Lambda function trigger with a failure condition is not a valid option, as triggers are used to invoke Lambda functions, not to send events from Lambda functions. Creating a status check alarm on the initial Lambda function will incur additional costs and complexity, and it will not capture the details of the failed
invocations. References
? Using AWS Lambda destinations
? Asynchronous invocation - AWS Lambda
? AWS Lambda Destinations: What They Are and Why to Use Them
? AWS Lambda Destinations: A Complete Guide | Dashbird

NEW QUESTION 41
A company's developer has deployed an application in AWS by using AWS CloudFormation The CloudFormation stack includes parameters in AWS Systems Manager Parameter Store that the application uses as configuration settings. The application can modify the parameter values
When the developer updated the stack to create additional resources with tags, the developer noted that the parameter values were reset and that the values ignored the latest changes made by the application. The developer needs to change the way the company deploys the CloudFormation stack. The developer also needs to avoid resetting the parameter values outside the stack.
Which solution will meet these requirements with the LEAST development effort?

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

Answer: D

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

NEW QUESTION 44
A developer is using AWS Amplify Hosting to build and deploy an application. The developer is receiving an increased number of bug reports from users. The developer wants to add end-to-end testing to the application to eliminate as many bugs as possible before the bugs reach production.
Which solution should the developer implement to meet these requirements?

A. Run the amplify add test command in the Amplify CLI.
B. Create unit tests in the applicatio
C. Deploy the unit tests by using the amplify push command in the Amplify CLI.
D. Add a test phase to the amplify.yml build settings for the application.
E. Add a test phase to the aws-exports.js file for the application.

Answer: C

Explanation:
The solution that will meet the requirements is to add a test phase to the amplify.yml build settings for the application. This way, the developer can run end-to-end tests on every code commit and catch any bugs before deploying to production. The other options either do not support end-to-end testing, or do not run tests automatically.
Reference: End-to-end testing

NEW QUESTION 48
A company has an ecommerce application. To track product reviews, the company's development team uses an Amazon DynamoDB table.
Every record includes the following
• A Review ID a 16-digrt universally unique identifier (UUID)
• A Product ID and User ID 16 digit UUlDs that reference other tables
• A Product Rating on a scale of 1-5
• An optional comment from the user
The table partition key is the Review ID. The most performed query against the table is to find the 10 reviews with the highest rating for a given product.
Which index will provide the FASTEST response for this query"?

A. A global secondary index (GSl) with Product ID as the partition key and Product Rating as the sort key
B. A global secondary index (GSl) with Product ID as the partition key and Review ID as the sort key
C. A local secondary index (LSI) with Product ID as the partition key and Product Rating as the sort key
D. A local secondary index (LSI) with Review ID as the partition key and Product ID as the sort key

Answer: A

Explanation:
This solution allows the fastest response for the query because it enables the query to use a single partition key value (the Product ID) and a range of sort key values (the Product Rating) to find the matching items. A global secondary index (GSI) is an index that has a partition key and an optional sort key that are different from those on the base table. A GSI can be created at any time and can be queried or scanned independently of the base table. A local secondary index (LSI) is an index that has the same partition key as the base table, but a different sort key. An LSI can only be created when the base table is created and must be queried together with the base table partition key. Using a GSI with Product ID as the partition key and Review ID as the sort key will not allow the query to use a range of sort key values to find the highest ratings. Using an LSI with Product ID as the partition key and Product Rating as the sort key will not work because Product ID is not the partition key of the base table. Using an LSI with Review ID as the partition key and Product ID as the sort key will not allow the query to use a single partition key value to find the matching items.
Reference: [Global Secondary Indexes], [Querying]

NEW QUESTION 49
A developer uses AWS CloudFormation to deploy an Amazon API Gateway API and an AWS Step Functions state machine The state machine must reference the API Gateway API after the CloudFormation template is deployed The developer needs a solution that uses the state machine to reference the API Gateway endpoint.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: A

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

? amazon-web-services aws-api-gateway terraform aws-step-functions

NEW QUESTION 50
A company is migrating legacy internal applications to AWS. Leadership wants to rewrite the internal employee directory to use native AWS services. A developer needs to create a solution for storing employee contact details and high-resolution photos for use with the new application.
Which solution will enable the search and retrieval of each employee's individual details and high-resolution photos using AWS APIs?

A. Encode each employee's contact information and photos using Base64. Store the information in an Amazon DynamoDB table using a sort key.
B. Store each employee's contact information in an Amazon DynamoDB table along with the object keys for the photos stored in Amazon S3.
C. Use Amazon Cognito user pools to implement the employee directory in a fully managed software-as-a-service (SaaS) method.
D. Store employee contact information in an Amazon RDS DB instance with the photos stored in Amazon Elastic File System (Amazon EFS).

Answer: B

Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance with seamless scalability. The developer can store each employee's contact information in a DynamoDB table along with the object keys for the photos stored in Amazon S3. Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. The developer can use AWS APIs to search and retrieve the employee details and photos from DynamoDB and S3.
References:
? [Amazon DynamoDB]
? [Amazon Simple Storage Service (S3)]

NEW QUESTION 53
A company has deployed infrastructure on AWS. A development team wants to create an AWS Lambda function that will retrieve data from an Amazon Aurora database. The Amazon Aurora database is in a private subnet in company's VPC. The VPC is named VPC1. The data is relational in nature. The Lambda function needs to access the data securely.
Which solution will meet these requirements?

A. Create the Lambda functio
B. Configure VPC1 access for the functio
C. Attach a security group named SG1 to both the Lambda function and the databas
D. Configure the security group inbound and outbound rules to allow TCP traffic on Port 3306.
E. Create and launch a Lambda function in a new public subnet that is in a new VPC named VPC2. Create a peering connection between VPC1 and VPC2.
F. Create the Lambda functio
G. Configure VPC1 access for the functio
H. Assign a security group named SG1 to the Lambda functio
I. Assign a second security group named SG2 to the databas
J. Add an inbound rule to SG1 to allow TCP traffic from Port 3306.
K. Export the data from the Aurora database to Amazon S3. Create and launch a Lambda function in VPC1. Configure the Lambda function query the data from Amazon S3.

Answer: A

Explanation:
AWS Lambda is a service that lets you run code without provisioning or managing servers. Lambda functions can be configured to access resources in a VPC, such as an Aurora database, by specifying one or more subnets and security groups in the VPC settings of the function. A security group acts as a virtual firewall that controls inbound and outbound traffic for the resources in a VPC. To allow a Lambda function to communicate with an Aurora database, both resources need to be associated with the same security group, and the security group rules need to allow TCP traffic on Port 3306, which is the default port for MySQL databases.
Reference: [Configuring a Lambda function to access resources in a VPC]

NEW QUESTION 57
A company is planning to use AWS CodeDeploy to deploy an application to Amazon Elastic Container Service (Amazon ECS) During the deployment of a new version of the application, the company initially must expose only 10% of live traffic to the new version of the deployed application. Then, after 15 minutes elapse, the company must route all the remaining live traffic to the new version of the deployed application.
Which CodeDeploy predefined configuration will meet these requirements?

A. CodeDeployDefault ECSCanary10Percent15Minutes
B. CodeDeployDefault LambdaCanary10Percent5Minutes
C. CodeDeployDefault LambdaCanary10Percent15Minutes
D. CodeDeployDefault ECSLinear10PercentEvery1 Minutes

Answer: A

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

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

NEW QUESTION 60
A developer is creating an AWS Lambda function. The Lambda function needs an external library to connect to a third-party solution The external library is a collection of files with a total size of 100 MB The developer needs to make the external library available to the Lambda execution environment and reduce the Lambda package space
Which solution will meet these requirements with the LEAST operational overhead?

A.

Create a Lambda layer to store the external library Configure the Lambda function to use the layer
B. Create an Amazon S3 bucket Upload the external library into the S3 bucke
C. Mount the S3 bucket folder in the Lambda function Import the library by using the proper folder in the mount point.
D. Load the external library to the Lambda function's /tmp directory during deployment of the Lambda packag
E. Import the library from the /tmp directory.
F. Create an Amazon Elastic File System (Amazon EFS) volum
G. Upload the external library to the EFS volume Mount the EFS volume in the Lambda functio
H. Import the library by using the proper folder in the mount point.

Answer: A

Explanation:
Create a Lambda layer to store the external library. Configure the Lambda function to use the layer. This will allow the developer to make the external library available to the Lambda execution environment without having to include it in the Lambda package, which will reduce the Lambda package space. Using a Lambda layer is a simple and straightforward solution that requires minimal operational overhead. https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html

NEW QUESTION 63
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 66
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 69
An application that is deployed to Amazon EC2 is using Amazon DynamoDB. The app cation calls the DynamoDB REST API Periodically the application receives a ProvisionedThroughputExceededException error when the application writes to a DynamoDB table.
Which solutions will mitigate this error MOST cost-effectively^ (Select TWO)

A. Modify the application code to perform exponential back off when the error is received.
B. Modify the application to use the AWS SDKs for DynamoDB.
C. Increase the read and write throughput of the DynamoDB table.
D. Create a DynamoDB Accelerator (DAX) cluster for the DynamoDB table.
E. Create a second DynamoDB table Distribute the reads and writes between the two tables.

Answer: AB

Explanation:
These solutions will mitigate the error most cost-effectively because they do not require increasing the provisioned throughput of the DynamoDB table or creating additional resources. Exponential backoff is a retry strategy that increases the waiting time between retries to reduce the number of requests sent to DynamoDB. The AWS SDKs for DynamoDB implement exponential backoff by default and also provide other features such as automatic pagination and encryption. Increasing the read and write throughput of the DynamoDB table, creating a DynamoDB Accelerator (DAX) cluster, or creating a second DynamoDB table will incur additional costs and complexity.
Reference: [Error Retries and Exponential Backoff in AWS], [Using the AWS SDKs with
DynamoDB]

NEW QUESTION 74
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 78
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.
I. Create a global variable that is outside the handler in the Lambda function to store the table name.

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 80
A developer at a company needs to create a small application mat makes the same API call once each flay at a designated time. The company does not have infrastructure in the AWS Cloud yet, but the company wants to implement this functionality on AWS.
Which solution meets these requirements in the MOST operationally efficient manner?

A. Use a Kubermetes cron job that runs on Amazon Elastic Kubemetes Sen/ice (Amazon EKS)
B. Use an Amazon Linux crontab scheduled job that runs on Amazon EC2
C. Use an AWS Lambda function that is invoked by an Amazon EventBridge scheduled event.
D. Use an AWS Batch job that is submitted to an AWS Batch job queue.

Answer: C

Explanation:
This solution meets the requirements in the most operationally efficient manner because it does not require any infrastructure provisioning or management. The developer can create a Lambda function that makes the API call and configure an EventBridge rule that triggers the function once a day at a designated time. This is a serverless solution that scales automatically and only charges for the execution time of the function.
Reference: [Using AWS Lambda with Amazon EventBridge], [Schedule Expressions for Rules]

NEW QUESTION 81
A company is planning to securely manage one-time fixed license keys in AWS. The company's development team needs to access the license keys in automaton scripts that run in Amazon EC2 instances and in AWS CloudFormation stacks.
Which solution will meet these requirements MOST cost-effectively?

A. Amazon S3 with encrypted files prefixed with "config"
B. AWS Secrets Manager secrets with a tag that is named SecretString
C. AWS Systems Manager Parameter Store SecureString parameters

D. CloudFormation NoEcho parameters

Answer: C

Explanation:
AWS Systems Manager Parameter Store is a service that provides secure, hierarchical storage for configuration data and secrets. Parameter Store supports SecureString parameters, which are encrypted using AWS Key Management Service (AWS KMS) keys. SecureString parameters can be used to store license keys in AWS and retrieve them securely from automation scripts that run in EC2 instances or CloudFormation stacks. Parameter Store is a cost-effective solution because it does not charge for storing parameters or API calls. Reference: Working with Systems Manager parameters

NEW QUESTION 85
A developer has created an AWS Lambda function that is written in Python. The Lambda function reads data from objects in Amazon S3 and writes data to an Amazon DynamoDB table. The function is successfully invoked from an S3 event notification when an object is created. However, the function fails when it attempts to write to the DynamoDB table.
What is the MOST likely cause of this issue?

A. The Lambda function's concurrency limit has been exceeded.
B. DynamoDB table requires a global secondary index (GSI) to support writes.
C. The Lambda function does not have IAM permissions to write to DynamoDB.
D. The DynamoDB table is not running in the same Availability Zone as the Lambda function.

Answer: C

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_lambda- access-dynamodb.html

NEW QUESTION 89
A company built a new application in the AWS Cloud. The company automated the bootstrapping of new resources with an Auto Scaling group by using AWS Cloudf-ormation templates. The bootstrap scripts contain sensitive data.
The company needs a solution that is integrated with CloudFormation to manage the sensitive data in the bootstrap scripts.
Which solution will meet these requirements in the MOST secure way?

A. Put the sensitive data into a CloudFormation paramete
B. Encrypt the CloudFormation templates by using an AWS Key Management Service (AWS KMS) key.
C. Put the sensitive data into an Amazon S3 bucket Update the CloudFormation templates to download the object from Amazon S3 during bootslrap.
D. Put the sensitive data into AWS Systems Manager Parameter Store as a secure string paramete
E. Update the CloudFormation templates to use dynamic references to specify template values.
F. Put the sensitive data into Amazon Elastic File System (Amazon EPS) Enforce EFS encryption after file system creatio
G. Update the CloudFormation templates to retrieve data from Amazon EFS.

Answer: C

Explanation:
This solution meets the requirements in the most secure way because it uses a service that is integrated with CloudFormation to manage sensitive data in encrypted form. AWS Systems Manager Parameter Store provides secure, hierarchical storage for configuration data management and secrets management. You can store sensitive data as secure string parameters, which are encrypted using an AWS Key Management Service (AWS KMS) key of your choice. You can also use dynamic references in your CloudFormation templates to specify template values that are stored in Parameter Store or Secrets Manager without having to include them in your templates. Dynamic references are resolved only during stack creation or update operations, which reduces exposure risks for sensitive data. Putting sensitive data into a CloudFormation parameter will not encrypt them or protect them from unauthorized access. Putting sensitive data into an Amazon S3 bucket or Amazon Elastic File System (Amazon EFS) will require additional configuration and integration with CloudFormation and may not provide fine-grained access control or encryption for sensitive data.
Reference: [What Is AWS Systems Manager Parameter Store?], [Using Dynamic References to Specify Template Values]

NEW QUESTION 90
A developer is creating an AWS Lambda function that searches for items from an Amazon DynamoDB table that contains customer contact information- The DynamoDB table items have the customer's email_address as the partition key and additional properties such as customer_type, name, and job_tltle.
The Lambda function runs whenever a user types a new character into the customer_type text input The developer wants the search to return partial matches of all the 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 as the partition key and email_address as the sort key Perform a query operation on the GSI by using the begvns_wth key condition expression With the email_address property
B. Add a global secondary index (GSI) to the DynamoDB table With ernail_address as the partition key and customer_type as the sort key Perform a query operation on the GSI by using the begins_wtth key condition expression With the email_address property.
C. Add a local secondary index (LSI) to the DynamoDB table With customer_type as the partition key and email_address as the sort key Perform a query operation on the LSI by using the begins_wlth key condition expression With the email_address property
D. Add a local secondary Index (LSI) to the DynamoDB table With job_tltle as the partition key and emad_address as the sort key Perform a query operation on the LSI by using the begins_wrth key condition expression With the email_address property

Answer: A

Explanation:
By adding a global secondary index (GSI) to the DynamoDB table with customer_type as the partition key and email_address as the sort key, the developer can perform a query operation on the GSI using the Begins_with key condition expression with the email_address property. This will return partial matches of all email_address properties of a specific customer_type.

NEW QUESTION 94
A developer accesses AWS CodeCommit over SSH. The SSH keys configured to access AWS CodeCommit are tied to a user with the following permissions:

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

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

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

B. "codecommit:Put*"

C. "codecommit:Update*"

D. "codecommit:*"

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

Answer: A

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

NEW QUESTION 99
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 103
.....

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