Exam Questions DVA-C02

DVA-C02

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

NEW QUESTION 1
A developer is deploying a company's application to Amazon EC2 instances The application generates gigabytes of data files each day The files are rarely accessed but the files must be available to the application's users within minutes of a request during the first year of storage The company must retain the files for 7 years.
How can the developer implement the application to meet these requirements MOST cost- effectively?

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

Answer: A

Explanation:
Amazon S3 Glacier Instant Retrieval is an archive storage class that delivers the lowest-cost storage for long-lived data that is rarely accessed and requires retrieval in
milliseconds. With S3 Glacier Instant Retrieval, you can save up to 68% on storage costs compared to using the S3 Standard-
Infrequent Access (S3 Standard-IA) storage class, when your data is accessed once per quarter. https://aws.amazon.com/s3/storage- classes/glacier/instant-retrieval/

NEW QUESTION 2
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 3
A developer maintains an Amazon API Gateway REST API. Customers use the API through a frontend UI and Amazon Cognito authentication.
The developer has a new version of the API that contains new endpoints and backward- incompatible interface changes. The developer needs to provide beta access to other developers on the team without affecting customers.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: A

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

NEW QUESTION 4
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 5
A developer maintains a critical business application that uses Amazon DynamoDB as the primary data store The DynamoDB table contains millions of documents and receives 30- 60 requests each minute The developer needs to perform processing in near-real time on the documents when they are added or updated in the DynamoDB table
How can the developer implement this feature with the LEAST amount of change to the existing application code?

A. Set up a cron job on an Amazon EC2 instance Run a script every hour to query the table for changes and process the documents
B. Enable a DynamoDB stream on the table Invoke an AWS Lambda function to process the documents.
C. Update the application to send a PutEvents request to Amazon EventBridg
D. Create an EventBridge rule to invoke an AWS Lambda function to process the documents.
E. Update the application to synchronously process the documents directly after the DynamoDB write

Answer: B

Explanation:
https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and- design-patterns/

NEW QUESTION 6
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 7
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 8
A company has an application that runs across multiple AWS Regions. The application is experiencing performance issues at irregular intervals. A developer must use AWS X-Ray to implement distributed tracing for the application to troubleshoot the root cause of the performance issues.
What should the developer do to meet this requirement?

A. Use the X-Ray console to add annotations for AWS services and user-defined services
B. Use Region annotation that X-Ray adds automatically for AWS services Add Region annotation for user-defined services
C. Use the X-Ray daemon to add annotations for AWS services and user-defined services
D. Use Region annotation that X-Ray adds automatically for user-defined services Configure X-Ray to add Region annotation for AWS services

Answer: B

Explanation:
 AWS X-Ray automatically adds Region annotation for AWS services that are integrated with X-Ray. This annotation indicates the AWS Region where the service is running. The developer can use this annotation to filter and group traces by Region and identify any performance issues related to cross-Region calls. The developer can also add Region annotation for user-defined services by using the X-Ray SDK. This option enables the developer to implement distributed tracing for the application that runs across multiple AWS Regions. References
? AWS X-Ray Annotations
? AWS X-Ray Concepts

NEW QUESTION 9
A company is using AWS CloudFormation to deploy a two-tier application. The application will use Amazon RDS as its backend database. The company wants a solution that will randomly generate the database password during deployment. The solution also must automatically rotate the database password without requiring changes to the application.
What is the MOST operationally efficient solution that meets these requirements'?

A. Use an AWS Lambda function as a CloudFormation custom resource to generate and rotate the password.
B. Use an AWS Systems Manager Parameter Store resource with the SecureString data type to generate and rotate the password.
C. Use a cron daemon on the application s host to generate and rotate the password.
D. Use an AWS Secrets Manager resource to generate and rotate the password.

Answer: D

Explanation:
 This solution will meet the requirements by using AWS Secrets Manager, which is a service that helps protect secrets such as database credentials by encrypting them with AWS Key Management Service (AWS KMS) and enabling automatic rotation of secrets. The developer can use an AWS Secrets Manager resource in AWS CloudFormation template, which enables creating and managing secrets as part of a CloudFormation stack. The developer can use an AWS::SecretsManager::Secret resource type to generate and rotate the password for accessing RDS database during deployment. The developer can also specify a RotationSchedule property for the secret resource, which defines how often to rotate the secret and which Lambda function to use for rotation logic. Option A is not optimal because it will use an AWS Lambda function as a CloudFormation custom resource, which may introduce additional complexity and overhead for creating and managing a custom resource and implementing rotation logic. Option B is not optimal because it will use an AWS Systems Manager Parameter Store resource with the SecureString data type, which does not support automatic rotation of secrets. Option C is not optimal because it will use a cron daemon on the application's host to generate and rotate the password, which may incur more costs and require more maintenance for running and securing a host.
References: [AWS Secrets Manager], [AWS::SecretsManager::Secret]

NEW QUESTION 10
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.
Store the credentials in AWS Systems Manager Parameter Store and enable rotatio
H. 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 10
A developer is troubleshooting an Amazon API Gateway API Clients are receiving HTTP 400 response errors when the clients try to access an endpoint of the API.
How can the developer determine the cause of these errors?

A. Create an Amazon Kinesis Data Firehose delivery stream to receive API call logs from API Gatewa

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

Answer: D

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

NEW QUESTION 13
A company 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 17
A developer is configuring an applications deployment environment in AWS CodePipeine. The application code is stored in a GitHub repository. The developer wants to ensure that the repository package's unit tests run in the new deployment environment. The deployment has already set the pipeline's source provider to GitHub and has specified the repository and branch to use in the deployment.
When combination of steps should the developer take next to meet these requirements with the least the LEAST overhead' (Select TWO).

A. Create an AWS CodeCommt projec
B. Add the repository package's build and test commands to the protects buildspec
C. Create an AWS CodeBuid projec
D. Add the repository package's build and test commands to the projects buildspec
E. Create an AWS CodeDeploy protec
F. Add the repository package's build and test commands to the project's buildspec
G. Add an action to the source stag
H. Specify the newly created project as the action provide
I. Specify the build attract as the actions input artifact.
J. Add a new stage to the pipeline alter the source stag
K. Add an action to the new stag
L. Speedy the newly created protect as the action provide
M. Specify the source artifact as the action's input artifact.

Answer: BE

Explanation:
This solution will ensure that the repository package's unit tests run in the new deployment environment with the least overhead because it uses AWS CodeBuild to build and test the code in a fully managed service, and AWS CodePipeline to orchestrate the deployment stages and actions. Option A is not optimal because it will use AWS CodeCommit instead of AWS CodeBuild, which is a source control service, not a build and test service. Option C is not optimal because it will use AWS CodeDeploy instead of AWS CodeBuild, which is a deployment service, not a build and test service. Option D is not optimal because it will add an action to

the source stage instead of creating a new stage, which will not follow the best practice of separating different deployment phases. References: AWS CodeBuild, AWS CodePipeline

NEW QUESTION 21
A developer wants to expand an application to run in multiple AWS Regions. The developer wants to copy Amazon Machine Images (AMIs) with the latest changes and create a new application stack in the destination Region. According to company requirements, all AMIs must be encrypted in all Regions. However, not all the AMIs that the company uses are encrypted.
How can the developer expand the application to run in the destination Region while meeting the encryption requirement?

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Amazon Machine Images (AMIs) are encrypted snapshots of EC2 instances that can be used to launch new instances. The developer can create new AMIs from the existing instances and specify encryption parameters. The developer can copy the encrypted AMIs to the destination Region and use them to create a new application stack. The developer can delete the unencrypted AMIs after the encryption process is complete. This solution will meet the encryption requirement and allow the developer to expand the application to run in the destination Region.
References:
? [Amazon Machine Images (AMI) - Amazon Elastic Compute Cloud]
? [Encrypting an Amazon EBS Snapshot - Amazon Elastic Compute Cloud]
? [Copying an AMI - Amazon Elastic Compute Cloud]

NEW QUESTION 25
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
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 DynamoDB 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 28
A developer is creating an application that will be deployed on IoT devices. The application will send data to a RESTful API that is deployed as an AWS Lambda function. The application will assign each API request a unique identifier. The volume of API requests from the application can randomly increase at any given time of day.
During periods of request throttling, the application might need to retry requests. The API must be able to handle duplicate requests without inconsistencies or data loss.
Which solution will meet these requirements?

A. Create an Amazon RDS for MySQL DB instanc
B. Store the unique identifier for each request in a database tabl
C. Modify the Lambda function to check the table for the identifier before processing the request.
D. Create an Amazon DynamoDB tabl
E. Store the unique identifier for each request in the tabl
F. Modify the Lambda function to check the table for the identifier before processing the request.

G. Create an Amazon DynamoDB tabl
H. Store the unique identifier for each request in the tabl

I. Modify the Lambda function to return a client error response when the function   receives a duplicate request.
J. Create an Amazon ElastiCache for Memcached instanc
K. Store the unique identifier for each request in the cach
L. Modify the Lambda function to check the cache for the identifier before processing the request.

Answer: B

Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that can store and retrieve any amount of data with high availability and performance. DynamoDB can handle concurrent requests from multiple IoT devices without throttling or data loss. To prevent duplicate requests from causing inconsistencies or data loss, the Lambda function can use DynamoDB conditional writes to check if the unique identifier for each request already exists in the table before processing the request. If the identifier exists, the function can skip or abort the request; otherwise, it can process the request and store the identifier in the table. Reference: Using conditional writes

NEW QUESTION 30
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 35
A company wants to share information with a third party. The third party has an HTTP API endpoint that the company can use to share the information. The company has the required API key to access the HTTP API.
The company needs a way to manage the API key by using code. The integration of the API key with the application code cannot affect application performance. Which solution will meet these requirements MOST securely?

A. Mastered
B. Not Mastered

Answer: A

Explanation:
AWS Secrets Manager is a service that helps securely store, rotate, and manage secrets such as API keys, passwords, and tokens. The developer can store the API credentials in AWS Secrets Manager and retrieve them at runtime by using the AWS SDK. This solution will meet the requirements of security, code management, and performance. Storing the API credentials in a local code variable or an S3 object is not secure, as it exposes the credentials to unauthorized access or leakage. Storing the API credentials in a DynamoDB table is also not secure, as it requires additional encryption and access control measures. Moreover, retrieving the credentials from S3 or DynamoDB may affect application performance due to network latency.
References:
? [What Is AWS Secrets Manager? - AWS Secrets Manager]
? [Retrieving a Secret - AWS Secrets Manager]

NEW QUESTION 37
A developer is testing an application that invokes an AWS Lambda function asynchronously. During the testing phase the Lambda function fails to process after two retries.
How can the developer troubleshoot the failure?

A. Configure AWS CloudTrail logging to investigate the invocation failures.
B. Configure Dead Letter Queues by sending events to Amazon SQS for investigation.
C. Configure Amazon Simple Workflow Service to process any direct unprocessed events.
D. Configure AWS Config to process any direct unprocessed events.

Answer: B

Explanation:
This solution allows the developer to troubleshoot the failure by capturing unprocessed events in a queue for further analysis. Dead Letter Queues (DLQs) are queues that store messages that could not be processed by a service, such as Lambda, for various reasons, such as configuration errors, throttling limits, or permissions issues. The developer can configure DLQs for Lambda functions by sending events to either an Amazon Simple Queue Service (SQS) queue or an Amazon Simple Notification Service (SNS) topic. The developer can then inspect the messages in the queue or topic to identify and fix the root cause of the failure. Configuring AWS CloudTrail logging will not capture invocation failures for asynchronous Lambda invocations, but only record API calls made by or on behalf of Lambda. Configuring Amazon Simple Workflow Service (SWF) or AWS Config will not process any direct unprocessed events, but require additional integration and configuration.
Reference: [Using AWS Lambda with DLQs], [Asynchronous invocation]

NEW QUESTION 38
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 39
A developer at a company needs to create a small application that makes the same API call once each day at a designated time. The company does not have infrastructure in the AWS Cloud yet, but the company wants to implement this functionality on AWS.
Which solution meets these requirements in the MOST operationally efficient manner?

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

Answer: C

Explanation:
The correct answer is C. Use an AWS Lambda function that is invoked by an Amazon EventBridge scheduled event.
* C. Use an AWS Lambda function that is invoked by an Amazon EventBridge scheduled event. This is correct. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda runs your code on a high-availability compute infrastructure and performs all of the administration of the compute resources, including server and operating system maintenance, capacity provisioning and automatic scaling, and logging1. Amazon EventBridge is a serverless event bus service that enables you to connect your applications with data from a variety of sources2. EventBridge can create rules that run on a schedule, either at regular intervals or at specific times and dates, and invoke targets such as Lambda functions3. This solution meets the requirements of creating a small application that makes the same API call once each day at a designated time, without requiring any infrastructure in the AWS Cloud or any operational overhead.
* A. Use a Kubernetes cron job that runs on Amazon Elastic Kubernetes Service (Amazon EKS). This is incorrect. Amazon EKS is a fully managed Kubernetes service that allows you to run containerized applications on AWS4. Kubernetes cron jobs are tasks that run periodically on a given schedule5. This solution could meet the functional requirements of creating a small application that makes the same API call once each day at a designated time, but it would not be the most operationally efficient manner. The company would need to provision and manage an EKS cluster, which would incur additional costs and complexity.
* B. Use an Amazon Linux crontab scheduled job that runs on Amazon EC2. This is incorrect. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud6. Crontab is a Linux utility that allows you to schedule commands or scripts to run automatically at a specified time or date7. This solution could meet the functional requirements of creating a small application that makes the same API call once each day at a designated time, but it would not be the most operationally efficient manner. The company would need to provision and manage an EC2 instance, which would incur additional costs and complexity.
* D. Use an AWS Batch job that is submitted to an AWS Batch job queue. This is incorrect. AWS Batch enables you to run batch computing workloads on the AWS Cloud8. Batch jobs are units of work that can be submitted to job queues, where they are executed in parallel or sequentially on compute environments9. This solution could meet the functional requirements of creating a small application that makes the same API call once each day at a designated time, but it would not be the most operationally efficient manner. The company would need to configure and manage an AWS Batch environment, which would incur additional costs and complexity.
References:
? 1: What is AWS Lambda? - AWS Lambda
? 2: What is Amazon EventBridge? - Amazon EventBridge
? 3: Creating an Amazon EventBridge rule that runs on a schedule - Amazon EventBridge
? 4: What is Amazon EKS? - Amazon EKS
? 5: CronJob - Kubernetes
? 6: What is Amazon EC2? - Amazon EC2
? 7: Crontab in Linux with 20 Useful Examples to Schedule Jobs - Tecmint
? 8: What is AWS Batch? - AWS Batch
? 9: Jobs - AWS Batch

NEW QUESTION 44
A company has an application that stores data in Amazon RDS instances. The application periodically experiences surges of high traffic that cause performance problems.
During periods of peak traffic, a developer notices a reduction in query speed in all database queries.

The team's technical lead determines that a multi-threaded and scalable caching solution should be used to offload the heavy read traffic. The solution needs to improve performance.

Which solution will meet these requirements with the LEAST complexity?

A. Use Amazon ElastiCache for Memcached to offload read requests from the main database.
B. Replicate the data to Amazon DynamoD
C. Set up a DynamoDB Accelerator (DAX) cluster.
D. Configure the Amazon RDS instances to use Multi-AZ deployment with one standby instanc
E. Offload read requests from the main database to the standby instance.
F. Use Amazon ElastiCache for Redis to offload read requests from the main database.

Answer: A

Explanation:

? Amazon ElastiCache for Memcached is a fully managed, multithreaded, and scalable in-memory key-value store that can be used to cache frequently accessed data and improve application performance1. By using Amazon ElastiCache for Memcached, the developer can reduce the load on the main database and handle high traffic surges more efficiently.
? To use Amazon ElastiCache for Memcached, the developer needs to create a cache cluster with one or more nodes, and configure the application to store and retrieve data from the cache cluster2. The developer can use any of the supported Memcached clients to interact with the cache cluster3. The developer can also use Auto Discovery to dynamically discover and connect to all cache nodes in a cluster4.
? Amazon ElastiCache for Memcached is compatible with the Memcached protocol, which means that the developer can use existing tools and libraries that work with

Memcached1. Amazon ElastiCache for Memcached also supports data partitioning, which allows the developer to distribute data among multiple nodes and scale out the cache cluster as needed.
? Using Amazon ElastiCache for Memcached is a simple and effective solution that meets the requirements with the least complexity. The developer does not need to change the database schema, migrate data to a different service, or use a different caching model. The developer can leverage the existing Memcached ecosystem and easily integrate it with the application.

NEW QUESTION 48
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 53
A developer is writing an application that will retrieve sensitive data from a third-party system. The application will format the data into a PDF file. The PDF file could be more than 1 MB. The application will encrypt the data to disk by using AWS Key Management Service (AWS KMS). The application will decrypt the file when a user requests to download it. The retrieval and formatting portions of the application are complete.
The developer needs to use the GenerateDataKey API to encrypt the PDF file so that the PDF file can be decrypted later. The developer needs to use an AWS KMS symmetric customer managed key for encryption.
Which solutions will meet these requirements?

A. Write the encrypted key from the GenerateDataKey API to disk for later us
B. Use the plaintext key from the GenerateDataKey API and a symmetric encryption algorithm to encrypt the file.
C. Write the plain text key from the GenerateDataKey API to disk for later us
D. Use the encrypted key from the GenerateDataKey API and a symmetric encryption algorithm to encrypt the file.
E. Write the encrypted key from the GenerateDataKey API to disk for later us
F. Use the plaintext key from the GenerateDataKey API to encrypt the file by using the KMS Encrypt API
G. Write the plain text key from the GenerateDataKey API to disk for later us
H. Use the encrypted key from the GenerateDataKey API to encrypt the file by using the KMS Encrypt API

Answer: A

Explanation:

? The GenerateDataKey API returns a data key that is encrypted under a symmetric encryption KMS key that you specify, and a plaintext copy of the same data key1. The data key is a random byte string that can be used with any standard encryption algorithm, such as AES or SM42. The plaintext data key can be used to encrypt or decrypt data outside of AWS KMS, while the encrypted data key can be stored with the encrypted data and later decrypted by AWS KMS1.
? In this scenario, the developer needs to use the GenerateDataKey API to encrypt
the PDF file so that it can be decrypted later. The developer also needs to use an AWS KMS symmetric customer managed key for encryption. To achieve this, the developer can follow these steps:

NEW QUESTION 57

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 62
A company's developer has deployed an application in AWS by using AWS CloudFormation The CloudFormation stack includes parameters in AWS Systems Manager Parameter Store that the application uses as configuration settings. The application can modify the parameter values
When the developer updated the stack to create additional resources with tags, the developer noted that the parameter values were reset and that the values ignored the latest changes made by the application. The developer needs to change the way the company deploys the CloudFormation stack. The developer also needs to avoid resetting the parameter values outside the stack.
Which solution will meet these requirements with the LEAST development effort?

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

Answer: D

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

NEW QUESTION 66
A developer 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 67
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 UUIDs 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 (GSI) with Product ID as the partition key and Product Rating as the sort key
B. A global secondary index (GSI) 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 71

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 73

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

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

Answer: C

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

NEW QUESTION 77

A developer is creating an AWS Lambda function in VPC mode An Amazon S3 event will invoke the Lambda function when an object is uploaded into an S3 bucket The Lambda function will process the object and produce some analytic results that will be recorded into a file Each processed object will also generate a log entry that will be recorded into a file.
Other Lambda functions. AWS services, and on-premises resources must have access to the result files and log file. Each log entry must also be appended to the same shared log file. The developer needs a solution that can share files and append results into an existing file.
Which solution should the developer use to meet these requirements?

A. Create an Amazon Elastic File System (Amazon EFS) file syste
B. Mount the EFS file system in Lambd
C. Store the result files and log file in the mount poin
D. Append the log entries to the log file.
E. Create an Amazon Elastic Block Store (Amazon EBS) Multi-Attach enabled volume Attach the EBS volume to all Lambda function
download the log file, append the log entries, and upload the modified log file to Amazon EBS
F. Update the Lambda function code to
G. Create a reference to the /tmp local director
H. Store the result files and log file by using the directory referenc
I. Append the log entry to the log file.
J. Create a reference to the /opt storage directory Store the result files and log file by using the directory reference Append the log entry to the log file

Answer: A

Explanation:
https://aws.amazon.com/blogs/compute/using-amazon-efs-for-aws-lambda-in-your-serverless-applications/

NEW QUESTION 82

A company has built an AWS Lambda function to convert large image files into output files that can be used in a third-party viewer application The company recently added a new module to the function to improve the output of the generated files However, the new module has increased the bundle size and has increased the time that is needed to deploy changes to the function code.

How can a developer increase the speed of the Lambda function deployment?

A. Use AWS CodeDeploy to deploy the function code
B. Use Lambda layers to package and load dependencies.
C. Increase the memory size of the function.
D. Use Amazon S3 to host the function dependencies

Answer: B

Explanation:
Using Lambda layers is a way to reduce the size of the deployment package and speed up the deployment process. Lambda layers are reusable components that can contain libraries, custom runtimes, or other dependencies. By using layers, the developer can separate the core function logic from the dependencies, and avoid uploading them every time the function code changes. Layers can also be shared across multiple functions or accounts, which can improve consistency and maintainability. References
? Working with AWS Lambda layers
? AWS Lambda Layers Best Practices
? Best practices for working with AWS Lambda functions

NEW QUESTION 87

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 92

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 94
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 97

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

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

Answer: C

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

NEW QUESTION 99

A developer wants to store information about movies. Each movie has a title, release year, and genre. The movie information also can include additional properties about the cast and production crew. This additional information is inconsistent across movies. For example, one movie might have an assistant director, and another movie might have an animal trainer.
The developer needs to implement a solution to support the following use cases:
For a given title and release year, get all details about the movie that has that title and release year.
For a given title, get all details about all movies that have that title. For a given genre, get all details about all movies in that genre. Which data store configuration will meet these requirements?

A. Create an Amazon DynamoDB tabl
B. Configure the table with a primary key that consists of the title as the partition key and the release year as the sort ke
C. Create a global secondary index that uses the genre as the partition key and the title as the sort key.

D. Create an Amazon DynamoDB tabl
E. Configure the table with a primary key that consists of the genre as the partition key and the release year as the sort ke
F. Create a global secondary index that uses the title as the partition key.
G. On an Amazon RDS DB instance, create a table that contains columns for title, release year, and genr
H. Configure the title as the primary key.
I. On an Amazon RDS DB instance, create a table where the primary key is the title and all other data is encoded into JSON format as one additional column.

Answer: A

Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance with seamless scalability. The developer can create a DynamoDB table and configure the table with a primary key that consists of the title as the partition key and the release year as the sort key. This will enable querying for a given title and release year efficiently. The developer can also create a global secondary index that uses the genre as the partition key and the title as the sort key. This will enable querying for a given genre efficiently. The developer can store additional properties about the cast and production crew as attributes in the DynamoDB table. These attributes can have different data types and structures, and they do not need to be consistent across items.
References:
? [Amazon DynamoDB]
? [Working with Queries - Amazon DynamoDB]
? [Working with Global Secondary Indexes - Amazon DynamoDB]

NEW QUESTION 104
A company is using Amazon API Gateway to invoke a new AWS Lambda function The company has Lambda function versions in its PROD and DEV environments. In each environment, there is a Lambda function alias pointing to the corresponding Lambda function version API Gateway has one stage that is configured to point at the PROD alias
The company wants to configure API Gateway to enable the PROD and DEV Lambda function versions to be simultaneously and distinctly available
Which solution will meet these requirements?

A. Enable a Lambda authorizer for the Lambda function alias in API Gateway Republish PROD and create a new stage for DEV Create API Gateway stage variables for the PROD and DEV stage
B. Point each stage variable to the PROD Lambda authorizer to the DEV Lambda authorizer.
C. Set up a gateway response in API Gateway for the Lambda function alia
D. Republish PROD and create a new stage for DE
E. Create gateway responses in API Gateway for PROD and DEV Lambda aliases
F. Use an environment variable for the Lambda function alias in API Gatewa
G. Republish PROD and create a new stage for developmen
H. Create API gateway environment variables for PROD and DEV stage
I. Point each stage variable to the PROD Lambda function alias to the DEV Lambda function alias.
J. Use an API Gateway stage variable to configure the Lambda function alias Republish PROD and create a new stage for development Create API Gateway stage variables for PROD and DEV stages Point each stage variable to the PROD Lambda function alias and to the DEV Lambda function alias

Answer: D

Explanation:
The best solution is to use an API Gateway stage variable to configure the Lambda function alias. This allows you to specify the Lambda function name and its alias or version using the syntax function_name:$ {stageVariables.variable_name} in the Integration Request. You can then create different stages in API Gateway, such as PROD and DEV, and assign different values to the stage variable for each stage. This way, you can invoke different Lambda function versions or aliases based on the stage that you are using, without changing the function name in the Integration Request. References
? Using API Gateway stage variables to manage Lambda functions
? How to point AWS API gateway stage to specific lambda function alias?
? Setting stage variables using the Amazon API Gateway console
? Amazon API Gateway stage variables reference

NEW QUESTION 107

A developer is modifying an existing AWS Lambda function White checking the code the developer notices hardcoded parameter various for an Amazon RDS for SQL Server user name password database host and port. There also are hardcoded parameter values for an Amazon DynamoOB table. an Amazon S3 bucket, and an Amazon Simple Notification Service (Amazon SNS) topic.

The developer wants to securely store the parameter values outside the code m an encrypted format and wants to turn on rotation for the credentials. The developer also wants to be able to reuse the parameter values from other applications and to update the parameter values without modifying code.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create an RDS database secret in AWS Secrets Manage
B. Set the user name password, database, host and por
C. Turn on secret rotatio
D. Create encrypted Lambda environment variables for the DynamoDB table, S3 bucket and SNS topic.
E. Create an RDS database secret in AWS Secrets Manage
F. Set the user name password, database, host and por
G. Turn on secret rotatio
H. Create Secure String parameters in AWS Systems Manager Parameter Store for the DynamoDB table, S3 bucket and SNS topic.
I. Create RDS database parameters in AWS Systems Manager Paramete
J. Store for the user name password, database, host and por
K. Create encrypted Lambda environment variables for me DynamoDB table, S3 bucket, and SNS topi
L. Create a Lambda function and set the logic for the credentials rotation task Schedule the credentials rotation task in Amazon EventBridge.
M. Create RDS database parameters in AWS Systems Manager Paramete
N. Store for the user name password database, host, and por
O. Store the DynamoDB tabl
P. S3 bucket, and SNS topic in Amazon S3 Create a Lambda function and set the logic for the credentials rotation Invoke the Lambda function on a schedule.

Answer: B

Explanation:
This solution will meet the requirements by using AWS Secrets Manager and AWS Systems Manager Parameter Store to securely store the parameter values outside the code in an encrypted format. AWS Secrets Manager is a service that helps protect secrets such as database credentials by encrypting them with AWS Key Management Service (AWS KMS) and enabling automatic rotation of secrets. The developer can create an RDS database secret in AWS Secrets Manager and set the user name, password, database, host, and port for accessing the RDS database. The developer can also turn on secret rotation, which will change the database credentials periodically according to a specified schedule or event. AWS Systems Manager Parameter Store is a service that provides secure and scalable storage for configuration data and secrets. The developer can create Secure String parameters in AWS Systems Manager Parameter Store for the DynamoDB table, S3 bucket, and SNS topic, which will encrypt them with AWS KMS. The developer can also reuse the parameter values from other applications and update them without modifying code. Option A is not optimal because it will create encrypted Lambda

environment variables for the

DynamoDB table, S3 bucket, and SNS topic, which may not be reusable or updatable without modifying code. Option C is not optimal because it will create RDS database parameters in AWS Systems Manager Parameter Store, which does not support automatic rotation of secrets. Option D is not optimal because it will store the DynamoDB table, S3 bucket, and SNS topic in Amazon S3, which may introduce additional costs and complexity for accessing configuration data.
References: AWS Secrets Manager, [AWS Systems Manager Parameter Store]

NEW QUESTION 109
A developer is designing an AWS Lambda function that creates temporary files that are less than 10 MB during invocation. The temporary files will be accessed and modified multiple times during invocation. The developer has no need to save or retrieve these files in the future.
Where should the temporary files be stored?

A. the /tmp directory
B. Amazon Elastic File System (Amazon EFS)
C. Amazon Elastic Block Store (Amazon EBS)
D. Amazon S3

Answer: A

Explanation:
 AWS Lambda is a service that lets developers run code without provisioning or managing servers. Lambda provides a local file system that can be used to store temporary files during invocation. The local file system is mounted under the /tmp directory and has a limit of 512 MB. The temporary files are accessible only by the Lambda function that created them and are deleted after the function execution ends. The developer can store temporary files that are less than 10 MB in the /tmp directory and access and modify them multiple times during invocation.
References:
? [What Is AWS Lambda? - AWS Lambda]
? [AWS Lambda Execution Environment - AWS Lambda]

NEW QUESTION 110
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 115
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 120
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 125
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

DVA-C02 Practice Exam Features:

* DVA-C02 Questions and Answers Updated Frequently

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

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

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