Amazon.DVA-C02.v2025-06-21.q138

Exam Code:	DVA-C02
Exam Name:	AWS Certified Developer - Associate
Certification Provider:	Amazon
Free Question Number:	138
Version:	v2025-06-21
# of views:	141
# of Questions views:	1380
https://www.dumpsdb.com/dumps/Amazon/DVA-C02/Amazon.DVA-C02.v2025-06-21.q138	

NEW QUESTION: 1
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 Canary10Percent10Minutes. Set the AutoPublishAlias property to the Lambda alias.
B. Set the Deployment Preference Type to LinearIOPercentEvery10Minutes. Set AutoPublishAlias property to the Lambda alias.
C. Set the Deployment Preference Type to CanaryIOPercentIOMinutes. Set the PreTraffic and PostTraffic properties to the Lambda alias.
D. Set the Deployment Preference Type to LinearIOPercentEveryIOMinutes. Set PreTraffic and Post Traffic properties to the Lambda alias.
Answer: (SHOW ANSWER)
The AWS Serverless Application Model (AWS SAM) comes built-in with CodeDeploy to provide gradual AWS Lambda deployments1. The DeploymentPreference property in AWS SAM allows you to specify the type of deployment that you want. The Canary10Percent10Minutes option means that 10 percent of your customer traffic is immediately shifted to your new version. After 10 minutes, all traffic is shifted to the new version1. The AutoPublishAlias property in AWS SAM allows AWS SAM to automatically create an alias that points to the updated version of the Lambda function1. Therefore, option A is correct.

NEW QUESTION: 2
An 1AM role is attached to an Amazon EC2 instance that explicitly denies access to all Amazon S3 API actions. The EC2 instance credentials file specifies the 1AM access key and secret access key, which allow full administrative access.
Given that multiple modes of 1AM access are present for this EC2 instance, which of the following is correct?
A. The EC2 instance will not be able to perform any S3 action on any S3 bucket.
B. The EC2 instance will only be able to list the S3 buckets.
C. The EC2 instance will be able to perform all actions on any S3 bucket.
D. The EC2 instance will only be able to list the contents of one S3 bucket at a time.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 3
A developer runs an application that displays scores for sports games on Amazon EC2 instances. The application uses a Redis client to retrieve the scores from an Amazon ElastiCache (Redis OSS) cluster.
The developer observes increased latency during operations on the cache because of connection failures to the cluster. The developer needs to resolve the latency issues.
A. Configure the Redis client to use an exponential backoff retry strategy to establish cache connections.
B. Store the scores in the application's memory. Perform bulk set operations on the scores that are stored in memory.
C. Configure the Redis client in the application to persist connections to the cluster by implementing a connection pool.
D. Deploy more nodes in the ElastiCache cluster. Update the Redis client to discover the new nodes.
Answer: (SHOW ANSWER)
Comprehensive Detailed Explanation with all AWS Reference
Why Option C is Correct:
Implementing a connection pool in the Redis client reduces connection overhead and avoids frequent connection establishment, which helps to reduce latency and connection failures.
Why Other Options are Incorrect:
Option A: Exponential backoff retry strategies help with transient failures but do not resolve latency caused by frequent connection establishment.
Option B: Storing scores in application memory adds complexity and risks inconsistency.
Option D: Adding more nodes is unnecessary unless the cluster is under heavy load. Latency due to connection failures is better addressed at the application level.
AWS Documentation Reference:
Amazon ElastiCache Best Practices

NEW QUESTION: 4

A company has an application that consists of different microservices that run inside an AWS account. The microservices are running in containers inside a single VPC. The number of microservices is constantly increasing. A developer must create a central logging solution for application logs.

A. Create a different Amazon CloudWatch Logs stream for each microservice.

B. Create an AWS CloudTrail trail to log all the API calls.

C. Configure VPC Flow Logs to track the communications between the microservices.

D. Use AWS Cloud Map to map the interactions of the microservices.

Answer: A (LEAVE A REPLY)

To create a central logging solution for microservices, using Amazon CloudWatch Logs is a recommended and effective approach. Here's why:

Amazon CloudWatch Logs Streams allow you to centralize logs from different services, which is crucial as the number of microservices increases.

Each microservice can have its own dedicated log stream within Amazon CloudWatch Logs, providing clear segregation of logs while still allowing centralized management.

This setup enables developers to monitor, search, and analyze logs efficiently using tools like CloudWatch Insights.

Other options like CloudTrail (B) are designed for API activity monitoring, not application logs. VPC Flow Logs (C) focus on network traffic rather than application behavior. AWS Cloud Map (D) is for service discovery and routing, not logging.

Reference:

AWS CloudWatch Logs documentation

Reference:

AWS CloudWatch Logs documentation

NEW QUESTION: 5

A developer is monitoring an application that runs on an Amazon EC2 Instance. The developer has configured a custom Amazon CloudWatch metric with data granularity of 1 second. It any issues occur, the developer wants to be notified within 30 seconds by Amazon Simple Notification Service (Amazon SNS).

What should the developer do to meet this requirement?

A. Configure a high-resolution CloudWatch alarm.

B. Use Amazon CloudWatch Logs Insights.

C. Set up a custom CloudWatch dashboard.

D. Change to a default CloudWatch metric.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 6

A developer is writing a web application that must share secure documents with end users. The documents are stored in a private Amazon S3 bucket. The application must allow only

authenticated users to download specific documents when requested, and only for a duration of 15 minutes.
How can the developer meet these requirements?
A. Modify the S3 bucket policy to only allow specific users to download the documents. Revert the change after 15 minutes.
B. Use server-side encryption with AWS KMS managed keys (SSE-KMS) and download the documents using HTTPS.
C. Create a presigned S3 URL using the AWS SDK with an expiration time of 15 minutes.
D. Copy the documents to a separate S3 bucket that has a lifecycle policy for deletion after 15 minutes.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 7
A developer is designing a fault-tolerant environment where client sessions will be saved.
How can the developer ensure that no sessions are lost if an Amazon EC2 instance fails?
A. Use sticky sessions with an Elastic Load Balancer target group.
B. Use Amazon DynamoDB to perform scalable session handling.
C. Use Elastic Load Balancer connection draining to stop sending requests to failing instances.
D. Use Amazon SOS to save session data.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 8
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 stages. Point each stage variable to the PROD Lambda authorizer to the DEV Lambda authorizer.
B. Set up a gateway response in API Gateway for the Lambda function alias. Republish PROD and create a new stage for DEV. Create gateway responses in API Gateway for PROD and DEV Lambda aliases
C. Use an environment variable for the Lambda function alias in API Gateway. Republish PROD and create a new stage for development. Create API gateway environment variables for PROD and DEV stages. Point each stage variable to the PROD Lambda function alias to the DEV Lambda function alias.
D. 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 (LEAVE A REPLY)

API Gateway Stages: Stages in API Gateway represent distinct environments (like PROD and DEV) allowing different configurations.

Stage Variables: Stage variables store environment-specific information, including Lambda function aliases.

Ease of Management: This solution offers a straightforward way to manage different Lambda function versions across environments.

Reference:

API Gateway Stages: https://docs.aws.amazon.com/apigateway/latest/developerguide/set-up-stages.html API Gateway Stage Variables: https://docs.aws.amazon.com/apigateway/latest/developerguide/stage-variables.html

NEW QUESTION: 9

A developer is creating a new batch application that will run on an Amazon EC2 instance. The application requires read access to an Amazon S3 bucket. The developer needs to follow security best practices to grant S3 read access to the application.

Which solution meets these requirements?

A. Add the permissions to an 1AM policy. Attach the policy to a user. Attach the user to the EC2 instance profile.

B. Add the permissions to an 1AM policy. Use 1AM web identity federation to access the S3 bucket with the policy.

C. Add the permissions to an 1AM policy. Attach the policy to a role. Attach the role to the EC2 instance profile.

D. Add the permissions inline to an 1AM group. Attach the group to the EC2 instance profile.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 10

A gaming application stores scores for players in an Amazon DynamoDB table that has four attributes: user_id, user_name, user_score, and user_rank. The users are allowed to update their names only. A user is authenticated by web identity federation.

Which set of conditions should be added in the policy attached to the role for the dynamodb:PutItem API call?

A. "Condition": {
"ForAllValues:StringEquals": {
"dynamodb:LeadingKeys": ["${www.amazon.com:user_id}"],
"dynamodb:Attributes": ["user_name"]
}
}

B. "Condition": {

"ForAllValues:StringEquals": {
"dynamodb:LeadingKeys": ["${www.amazon.com:user_name}"],
"dynamodb:Attributes": ["user_id"]
}
}
C. "Condition": {
"ForAllValues:StringEquals": {
"dynamodb:LeadingKeys": ["${www.amazon.com:user_id}"],
"dynamodb:Attributes": ["user_name", "user_id"]
}
}
D. "Condition": {
"ForAllValues:StringEquals": {
"dynamodb:LeadingKeys": ["${www.amazon.com:user_name}"],
"dynamodb:Attributes": ["username", "userid"]
}
}
Answer: A (LEAVE A REPLY)
The correct policy condition ensures that:
The LeadingKeys condition restricts operations to the authenticated user's user_id.
The Attributes condition limits the updatable attributes to user_name.
Explanation of Choices:
Option A: Correctly enforces both the key restriction (dynamodb:LeadingKeys) and ensures only the user_name attribute can be updated.
Option B, C, D: Use incorrect conditions, such as referencing user_name in the LeadingKeys or including other attributes like user_id in updatable fields.
Reference:
AWS DynamoDB Condition Keys Documentation
Reference:
AWS DynamoDB Condition Keys Documentation

NEW QUESTION: 11
A developer needs to troubleshoot an AWS Lambda function in a development environment. The Lambda function is configured in VPC mode and needs to connect to an existing Amazon RDS for SOL Server DB instance. The DB instance is deployed in a private subnet and accepts connections by using port 1433.
When the developer tests the function, the function reports an error when it tries to connect to the database.
Which combination of steps should the developer take to diagnose this issue? (Select TWO.)

A. Check that the function's security group has outbound access on port 1433 to the DB instance's security group. Check that the DB instance's security group has inbound access on port 1433 from the function's security group.
B. Check that the function's security group has Inbound access on port 1433 from the DB Instance's security group. Check that the DB instance's security group has outbound access on port 1433 to the function's security group.
C. Check that the function's execution rote permissions include ec2: CreateNetworkInterface. ec2: DescribeNetworkInterfaces. and ec2: DeleteNetworkInterface.
D. Check that the function's execution role permissions include rds:DescribeDBInstances, rds: ModifyDB Instance, and rds:DescribeDBSecurityGroups for the DB instance.
E. Check that the VPC is set up for a NAT gateway. Check that the DB instance has the public access option turned on.
Answer: A,C (LEAVE A REPLY)

NEW QUESTION: 12
A developer is creating a template that uses AWS CloudFormation to deploy an application. The application is serverless and uses Amazon API Gateway, Amazon DynamoDB, and AWS Lambda.
Which AWS service or tool should the developer use to define serverless resources in YAML?
A. CloudFormation serverless intrinsic functions
B. AWS Elastic Beanstalk
C. AWS Serverless Application Model (AWS SAM)
D. AWS Cloud Development Kit (AWS CDK)
Answer: C (LEAVE A REPLY)
AWS Serverless Application Model (AWS SAM) is an open-source framework that enables developers to build and deploy serverless applications on AWS. AWS SAM uses a template specification that extends AWS CloudFormation to simplify the definition of serverless resources such as API Gateway, DynamoDB, and Lambda. The developer can use AWS SAM to define serverless resources in YAML and deploy them using the AWS SAM CLI.
Reference:
[What Is the AWS Serverless Application Model (AWS SAM)? - AWS Serverless Application Model]
[AWS SAM Template Specification - AWS Serverless Application Model]

NEW QUESTION: 13
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
Answer: A (LEAVE A REPLY)
the resource.
Explanation:
CloudFormation and Dynamic Reference: The DefinitionSubstitutions property in CloudFormation allows you to pass values into Step Functions state machines at runtime.
Cost-Effectiveness: This solution is cost-effective as it leverages CloudFormation's built-in capabilities, avoiding the need for additional services like Secrets Manager or AppConfig.
Reference:
AWS Step Functions State Machine:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-stepfunctions-statemachine.html CloudFormation DefinitionSubstitutions:
https://github.com/aws-cloudformation/aws-cloudformation-resource-providers-stepfunctions/issues/14

NEW QUESTION: 14
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 group. Configure the filter to send an SNS notification when the number of errors exceeds the defined threshold within a 5-minute period.
C. 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
D. 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 metric. Configure the alarm to send an SNS notification when the number of errors exceeds the defined threshold within a 5-minute period.

Answer: D (LEAVE A REPLY)
CloudWatch for Log Analysis: CloudWatch is the best fit here because logs are already centralized. Here's the process:
Metric Filter: Create a metric filter on the CloudWatch Logs log group. Design a pattern to specifically identify application error messages.
Custom Metric: This filter generates a new custom CloudWatch metric (e.g., ApplicationErrors). This metric tracks the error count.
CloudWatch Alarm: Create an alarm on the ApplicationErrors metric. Configure the alarm with your desired threshold and a 5-minute evaluation period.
SNS Action: Set the alarm to trigger an SNS notification when it enters the alarm state.
Reference:
CloudWatch Metric Filters:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/MonitoringLogData.html
CloudWatch Alarms:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/AlarmThatSendsEmail.html

NEW QUESTION: 15
A company is using an Amazon API Gateway REST API endpoint as a webhook to publish events from an on-premises source control management (SCM) system to Amazon EventBridge. The company has configured an EventBridge rule to listen for the events and to control application deployment in a central AWS account. The company needs to receive the same events across multiple receiver AWS accounts.
How can a developer meet these requirements without changing the configuration of the SCM system?
A. Grant permission to the central AWS account for EventBridge to access the receiver AWS accounts. Add an EventBridge event bus on the receiver AWS accounts as the targets to the existing EventBridge rule.
B. Convert the API Gateway type from REST API to HTTP API.
C. Deploy the API Gateway REST API to all the receiver AWS accounts. Create as many SCM webhooks as the number of AWS accounts.
D. Deploy the API Gateway REST API to all the required AWS accounts. Use the same custom domain name for all the gateway endpoints so that a single SCM webhook can be used for all events from all accounts.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 16
A company runs an application on AWS The application uses an AWS Lambda function that is configured with an Amazon Simple Queue Service (Amazon SQS) queue called high priority queue as the event source A developer is updating the Lambda function with another SQS queue called low priority queue as the event source The Lambda function must always read up

to 10 simultaneous messages from the high priority queue before processing messages from low priority queue. The Lambda function must be limited to 100 simultaneous invocations. Which solution will meet these requirements'?

A. Set the event source mapping batch size to 10 for the high priority queue and to 90 for the low priority queue

B. Set the delivery delay to 0 seconds for the high priority queue and to 10 seconds for the low priority queue

C. Set the event source mapping maximum concurrency to 10 for the high priority queue and to 90 for the low priority queue

D. Set the event source mapping batch window to 10 for the high priority queue and to 90 for the low priority queue

Answer: C (LEAVE A REPLY)

Lambda Concurrency: The 'maximum concurrency' setting in event source mappings controls the maximum number of simultaneous invocations Lambda allows for that specific source. Prioritizing Queues: Setting a lower maximum concurrency for the 'high priority queue' ensures it's processed first while allowing more concurrent invocations from the 'low priority queue'. Batching: Batch size settings affect the number of messages Lambda retrieves from a queue per invocation, which is less relevant to the prioritization requirement.

Reference:

Lambda Event Source Mappings: https://docs.aws.amazon.com/lambda/latest/dg/invocation-eventsourcemapping.html Lambda Concurrency: https://docs.aws.amazon.com/lambda/latest/dg/configuration-concurrency.html

Valid DVA-C02 Dumps shared by TrainingQuiz.com for Helping Passing DVA-C02 Exam! TrainingQuiz.com now offer the **newest DVA-C02 exam dumps**, the TrainingQuiz.com DVA-C02 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C02 dumps with Test Engine here: https://www.trainingquiz.com/DVA-C02-practice-quiz.html (**649** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

NEW QUESTION: 17

A developer is creating a mobile application that will not require users to log in.

What is the MOST efficient method to grant users access to AWS resources'?

A. Use an identity provider to securely authenticate with the application.

B. Create an AWS Lambda function to create an 1AM user when a user accesses the application.

C. Create credentials using AWS KMS and apply these credentials to users when using the application.

D. Use Amazon Cognito to associate unauthenticated users with an IAM role that has limited access to resources.

Answer: (SHOW ANSWER)

This solution is the most efficient method to grant users access to AWS resources without requiring them to log in. Amazon Cognito is a service that provides user sign-up, sign-in, and access control for web and mobile applications. Amazon Cognito identity pools support both authenticated and unauthenticated users. Unauthenticated users receive access to your AWS resources even if they aren't logged in with any of your identity providers (IdPs). You can use Amazon Cognito to associate unauthenticated users with an IAM role that has limited access to resources, such as Amazon S3 buckets or DynamoDB tables. This degree of access is useful to display content to users before they log in or to allow them to perform certain actions without signing up. Using an identity provider to securely authenticate with the application will require users to log in, which does not meet the requirement. Creating an AWS Lambda function to create an IAM user when a user accesses the application will incur unnecessary costs and complexity, and may pose security risks if not implemented properly. Creating credentials using AWS KMS and applying them to users when using the application will also incur unnecessary costs and complexity, and may not provide fine-grained access control for resources.

NEW QUESTION: 18

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 (LEAVE A REPLY)

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.

NEW QUESTION: 19

A software company is launching a multimedia application. The application will allow guest users to access sample content before the users decide if they want to create an account to gain full access. The company wants to implement an authentication process that can identify users who have already created an account. The company also needs to keep track of the number of guest users who eventually create an account.
Which combination of steps will meet these requirements? {Select TWO.)
A. Create an Amazon Cognito identity pool. Configure the identity pool to allow unauthenticated users. Exchange unique identity for temporary credentials that allow all users to assume a role.
B. Create an Amazon CloudFront distribution. Configure the distribution to allow unauthenticated users. Exchange user tokens for temporary credentials that allow all users to assume a role.
C. Create a role for authenticated users that allows access to all content. Create a role for unauthenticated users that allows access to only the sample content.
D. Allow all users to access the sample content by default. Create a role for authenticated users that allows access to the other content.
E. Create an Amazon Cognito user pool. Configure the user pool to allow unauthenticated users. Exchange user tokens for temporary credentials that allow authenticated users to assume a role.
Answer: A,C (LEAVE A REPLY)

NEW QUESTION: 20
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 parameter. Encrypt the CloudFormation templates by using an AWS Key Management Service (AWS KMS) key.
B. Put the sensitive data into an Amazon S3 bucket Update the CloudFormation templates to download the object from Amazon S3 during bootslrap.
C. Put the sensitive data into AWS Systems Manager Parameter Store as a secure string parameter. Update the CloudFormation templates to use dynamic references to specify template values.
D. Put the sensitive data into Amazon Elastic File System (Amazon EPS) Enforce EFS encryption after file system creation. Update the CloudFormation templates to retrieve data from Amazon EFS.
Answer: C (LEAVE A REPLY)
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.

NEW QUESTION: 21
A developer has written an AWS Lambda function. The function is CPU-bound. The developer wants to ensure that the function returns responses quickly.
How can the developer improve the function's performance?
A. Increase the function's CPU core count.
B. Increase the function's memory.
C. Increase the function's reserved concurrency.
D. Increase the function's timeout.
Answer: B (LEAVE A REPLY)
The amount of memory you allocate to your Lambda function also determines how much CPU and network bandwidth it gets. Increasing the memory size can improve the performance of CPU-bound functions by giving them more CPU power. The CPU allocation is proportional to the memory allocation, so a function with 1 GB of memory has twice the CPU power of a function with 512 MB of memory. Reference: AWS Lambda execution environment

NEW QUESTION: 22
A developer designed an application on an Amazon EC2 instance The application makes API requests to objects in an Amazon S3 bucket Which combination of steps will ensure that the application makes the API requests in the MOST secure manner? (Select TWO.)
A. Create an IAM user that has permissions to the S3 bucket. Add the user to an 1AM group
B. Create an IAM role that has permissions to the S3 bucket
C. Add the IAM role to an instance profile. Attach the instance profile to the EC2 instance.
D. Create an 1AM role that has permissions to the S3 bucket Assign the role to an 1AM group
E. Store the credentials of the IAM user in the environment variables on the EC2 instance
Answer: B,C (LEAVE A REPLY)
IAM Roles for EC2: IAM roles are the recommended way to provide AWS credentials to applications running on EC2 instances. Here's how this works:
You create an IAM role with the necessary permissions to access the target S3 bucket.
You create an instance profile and associate the IAM role with this profile.

When launching the EC2 instance, you attach this instance profile.
Temporary Security Credentials: When the application on the EC2 instance needs to access S3, it doesn't directly use access keys. Instead, the AWS SDK running on the instance retrieves temporary security credentials associated with the role. These are rotated automatically by AWS.
Reference:
IAM Roles for Amazon EC2:
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html
Temporary Security Credentials:
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp.html

NEW QUESTION: 23
A developer is using AWS CodeDeploy to launch an application onto Amazon EC2 instances. The application deployment fails during testing. The developer notices an IAM_ROLE_PERMISSIONS error code in Amazon CloudWatch logs.
What should the developer do to resolve the error?
A. Attach the AWSCodeDeployRole policy to the CodeDeploy service role.
B. Attach the AWSCodeDeployRoleECS policy to the CodeDeploy service role.
C. Ensure that the deployment group is using the correct role name for the CodeDeploy service role.
D. Ensure the CodeDeploy agent is installed and running on all instances in the deployment group.
Answer: (SHOW ANSWER)

NEW QUESTION: 24
A financial company must store original customer records for 10 years for legal reasons. A complete record contains personally identifiable information (PII). According to local regulations, PII is available to only certain people in the company and must not be shared with third parties. The company needs to make the records available to third-party organizations for statistical analysis without sharing the PII.
A developer wants to store the original immutable record in Amazon S3. Depending on who accesses the S3 document, the document should be returned as is or with all the PII removed. The developer has written an AWS Lambda function to remove the PII from the document. The function is named removePii.
What should the developer do so that the company can meet the PII requirements while maintaining only one copy of the document?
A. Set up an S3 event notification that invokes the removePii function when an S3 GET request is made. Call Amazon S3 by using a GET request to access the object without PII.
B. Set up an S3 event notification that invokes the removePii function when an S3 PUT request is made. Call Amazon S3 by using a PUT request to access the object without PII.

C. Create an S3 Object Lambda access point from the S3 console. Select the removePii function. Use S3 Access Points to access the object without PII.
D. Create an S3 access point from the S3 console. Use the access point name to call the GetObjectLegalHold S3 API function. Pass in the removePii function name to access the object without PII.
Answer: (SHOW ANSWER)
S3 Object Lambda allows you to add your own code to process data retrieved from S3 before returning it to an application. You can use an AWS Lambda function to modify the data, such as removing PII, redacting confidential information, or resizing images. You can create an S3 Object Lambda access point and associate it with your Lambda function. Then, you can use the access point to request objects from S3 and get the modified data back. This way, you can maintain only one copy of the original document in S3 and apply different transformations depending on who accesses it. Reference: Using AWS Lambda with Amazon S3

NEW QUESTION: 25
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 (LEAVE A REPLY)
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.

NEW QUESTION: 26
A developer is creating an ecommerce workflow in an AWS Step Functions state machine that includes a HTTP Task state. The task passes shipping information and order details to an endpoint.
The developer needs to test the workflow to confirm that the HTTP headers and body are correct and that the responses meet expectations.
A. Use the TestState API to invoke only the HTTP Task. Set the inspection level to TRACE.
B. Use the TestState API to invoke the state machine. Set the inspection level to DEBUG.
C. Use the data flow simulator to invoke only the HTTP Task. View the request and response data.
D. Change the log level of the state machine to ALL. Run the state machine.
Answer: (SHOW ANSWER)
Comprehensive and Detailed Step-by-Step
To confirm that the HTTP headers, body, and responses meet expectations, you need to test the specific HTTP Task state in isolation and inspect the details.
Option A: TestState API with TRACE:
The TestState API allows developers to test individual states in a state machine without executing the entire workflow.
Setting the inspection level to TRACE provides detailed information about the HTTP request and response, including headers, body, and status codes.
This option provides the precise and granular testing required to verify the HTTP Task functionality.
Why Other Options Are Incorrect:
Option B: The DEBUG inspection level provides less detailed information than TRACE and focuses on general debugging, not a detailed view of HTTP interactions.
Option C: Step Functions does not have a "data flow simulator" to test individual tasks; this option is not valid.
Option D: Changing the state machine's log level to ALL increases logging granularity for the entire state machine but does not allow isolated testing of a specific HTTP Task.
Reference:
AWS Step Functions: Testing State Machines

NEW QUESTION: 27
A developer is using an AWS CloudFormation template to create a pipeline in AWS CodePipeline. The template creates an Amazon S3 bucket that the pipeline references in a source stage. The template also creates an AWS CodeBuild project for a build stage. The pipeline sends notifications to an Amazon SNS topic. Logs for the CodeBuild project are stored in Amazon CloudWatch Logs.

The company needs to ensure that the pipeline's artifacts are encrypted with an existing customer-managed AWS KMS key. The developer has granted the pipeline permissions to use the KMS key.
Which additional step will meet these requirements?
A. Create an Amazon S3 gateway endpoint that the pipeline can access.
B. In the CloudFormation template, use the KMS key to encrypt the logs in CloudWatch Logs.
C. Apply an S3 bucket policy that ensures the pipeline sends only encrypted objects to the S3 bucket.
D. Configure the notification topic to use the existing KMS key to enable encryption with the existing KMS key.
Answer: C (LEAVE A REPLY)
Comprehensive Detailed Explanation with all AWS Reference
Why Option C is Correct:
Ensuring that pipeline artifacts are encrypted with a customer-managed AWS KMS key involves configuring the S3 bucket policy to require encryption. This policy ensures all objects uploaded to the bucket are encrypted with the specified KMS key.
Why Other Options are Incorrect:
Option A: A gateway endpoint improves S3 access efficiency but does not enforce encryption.
Option B: Encrypting CloudWatch Logs is unrelated to securing pipeline artifacts.
Option D: Configuring SNS for encryption does not affect the artifacts stored in the S3 bucket.
AWS Documentation Reference:
Using Server-Side Encryption with S3 Bucket Policies

NEW QUESTION: 28
A company has developed a new serverless application using AWS Lambda functions that will be deployed using the AWS Serverless Application Model (AWS SAM) CLI.
Which step should the developer complete prior to deploying the application?
A. Compress the application to a zip file and upload it into AWS Lambda.
B. Test the new AWS Lambda function by first tracing it m AWS X-Ray.
C. Bundle the serverless application using a SAM package.
D. Create the application environment using the eb create my-env command.
Answer: C (LEAVE A REPLY)
This step should be completed prior to deploying the application because it prepares the application artifacts for deployment. The AWS Serverless Application Model (AWS SAM) is a framework that simplifies building and deploying serverless applications on AWS. The AWS SAM CLI is a command-line tool that helps you create, test, and deploy serverless applications using AWS SAM templates. The sam package command bundles the application artifacts, such as Lambda function code and API definitions, and uploads them to an Amazon S3 bucket. The command also returns a CloudFormation template that is ready to be deployed with the sam deploy command. Compressing the application to a zip file and uploading it to AWS Lambda will not work because it does not use AWS SAM templates or CloudFormation.

Testing the new Lambda function by first tracing it in AWS X-Ray will not prepare the application for deployment, but only monitor its performance and errors. Creating the application environment using the eb create my-env command will not work because it is a command for AWS Elastic Beanstalk, not AWS SAM.

NEW QUESTION: 29
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 archive. Deploy each Lambda function with its own copy of the library.
B. Create a Lambda layer with the required Python library. Use the Lambda layer in both Lambda functions.
C. Combine the two Lambda functions into one Lambda function. Deploy the Lambda function as a single .zip file archive.
D. Download the Python library to an S3 bucket. Program the Lambda functions to reference the object URLs.
Answer: (SHOW ANSWER)
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.
Reference:
[What Is AWS Lambda? - AWS Lambda]
[AWS Lambda Layers - AWS Lambda]

NEW QUESTION: 30
A company had an Amazon RDS for MySQL DB instance that was named mysql-db. The DB instance was deleted within the past 90 days. A developer needs to find which 1AM user or role deleted the DB instance in the AWS environment. Which solution will provide this information?
A. Retrieve the AWS CloudTrail events for the resource mysql-db where the event name is DeleteDBInstance. Inspect each event.

B. Retrieve the Amazon CloudWatch log events from the most recent log stream within the rds/mysql-db log group. Inspect the log events.
C. Retrieve the AWS X-Ray trace summaries. Filter by services with the name mysql-db. Inspect the ErrorRootCauses values within each summary.
D. Retrieve the AWS Systems Manager deletions inventory Filter the inventory by deletions that have a TypeName value of RDS. Inspect the deletion details.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 31
A company wants to deploy and maintain static websites on AWS. Each website's source code is hosted in one of several version control systems, including AWS CodeCommit, Bitbucket, and GitHub.
The company wants to implement phased releases by using development, staging, user acceptance testing, and production environments in the AWS Cloud. Deployments to each environment must be started by code merges on the relevant Git branch. The company wants to use HTTPS for all data exchange. The company needs a solution that does not require servers to run continuously.
Which solution will meet these requirements with the LEAST operational overhead?
A. Host each website by using AWS Amplify with a serverless backend. Conned the repository branches that correspond to each of the desired environments. Start deployments by merging code changes to a desired branch.
B. Host each website in AWS Elastic Beanstalk with multiple environments. Use the EB CLI to link each repository branch. Integrate AWS CodePipeline to automate deployments from version control code merges.
C. Host each website in different Amazon S3 buckets for each environment. Configure AWS CodePipeline to pull source code from version control. Add an AWS CodeBuild stage to copy source code to Amazon S3.
D. Host each website on its own Amazon EC2 instance. Write a custom deployment script to bundle each website's static assets. Copy the assets to Amazon EC2. Set up a workflow to run the script when code is merged.
Answer: A (LEAVE A REPLY)
AWS Amplify is a set of tools and services that enables developers to build and deploy full-stack web and mobile applications that are powered by AWS. AWS Amplify supports hosting static websites on Amazon S3 and Amazon CloudFront, with HTTPS enabled by default. AWS Amplify also integrates with various version control systems, such as AWS CodeCommit, Bitbucket, and GitHub, and allows developers to connect different branches to different environments. AWS Amplify automatically builds and deploys the website whenever code changes are merged to a connected branch, enabling phased releases with minimal operational overhead. Reference: AWS Amplify Console

Valid DVA-C02 Dumps shared by TrainingQuiz.com for Helping Passing DVA-C02 Exam! TrainingQuiz.com now offer the **newest DVA-C02 exam dumps**, the TrainingQuiz.com DVA-C02 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C02 dumps with Test Engine here: https://www.trainingquiz.com/DVA-C02-practice-quiz.html (**649** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

NEW QUESTION: 32
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 use. Use the plaintext key from the GenerateDataKey API and a symmetric encryption algorithm to encrypt the file.
B. Write the plain text key from the GenerateDataKey API to disk for later use. Use the encrypted key from the GenerateDataKey API and a symmetric encryption algorithm to encrypt the file.
C. Write the encrypted key from the GenerateDataKey API to disk for later use. Use the plaintext key from the GenerateDataKey API to encrypt the file by using the KMS Encrypt API
D. Write the plain text key from the GenerateDataKey API to disk for later use. Use the encrypted key from the GenerateDataKey API to encrypt the file by using the KMS Encrypt API
Answer: (SHOW ANSWER)
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: Call the GenerateDataKey API with the symmetric customer managed key ID and the desired length or specification of the data key. The API will return an encrypted data key and a plaintext data key.

Write the encrypted data key to disk for later use. This will allow the developer to decrypt the data key and the PDF file later by using AWS KMS.
Use the plaintext data key and a symmetric encryption algorithm to encrypt the PDF file. The developer can use any standard encryption library or tool to perform this operation, such as OpenSSL or AWS Encryption SDK.
Discard the plaintext data key from memory as soon as possible after using it. This will prevent unauthorized access or leakage of the data key.

NEW QUESTION: 33
In a move toward using microservices, a company's management team has asked all development teams to build their services so that API requests depend only on that service's data store. One team is building a Payments service which has its own database; the service needs data that originates in the Accounts database. Both are using Amazon DynamoDB.
What approach will result in the simplest, decoupled, and reliable method to get near-real time updates from the Accounts database?
A. Use Amazon DynamoDB Streams to deliver all changes from the Accounts database to the Payments database.
B. Use Amazon ElastiCache in Payments, with the cache updated by triggers in the Accounts database.
C. Use AWS Glue to perform frequent ETL updates from the Accounts database to the Payments database.
D. Use Amazon Data Firehose to deliver all changes from the Accounts database to the Payments database.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 34
A developer is deploying an AWS Lambda function The developer wants the ability to return to older versions of the function quickly and seamlessly.
How can the developer achieve this goal with the LEAST operational overhead?
A. Use AWS OpsWorks to perform blue/green deployments.
B. Use a function alias with different versions.
C. Maintain deployment packages for older versions in Amazon S3.
D. Use AWS CodePipeline for deployments and rollbacks.
Answer: (SHOW ANSWER)
A function alias is a pointer to a specific Lambda function version. You can use aliases to create different environments for your function, such as development, testing, and production. You can also use aliases to perform blue/green deployments by shifting traffic between two versions of your function gradually. This way, you can easily roll back to a previous version if something goes wrong, without having to redeploy your code or change your configuration.
Reference: AWS Lambda function aliases

NEW QUESTION: 35
A company has a serverless application that uses Amazon API Gateway backed by AWS Lambda proxy integration. The company is developing several backend APIs. The company needs a landing page to provide an overview of navigation to the APIs.
A developer creates a new /LandingPage resource and a new GET method that uses mock integration.
What should the developer do next to meet these requirements?
A. Configure the integration request mapping template with Content-Type of text/html. In the integration request mapping template, include the LandingPage HTML code that references the APIs. Configure the integration response mapping template with Content-Type of application/json and statusCode of 200.
B. Configure the Integration request mapping template with Content-Type of application/json. In the integration request mapping template, include the LandingPage HMTL code that references the APIs. Configure the integration response mapping template with Content-Type of text/html and statusCode of 200.
C. Configure the integration request mapping template with Content-Type of text/html and statusCode of 200. Configure the integration response mapping template with Content-Type of application/json. In the integration response mapping template, include the LandingPage HTML code that references the APIs.
D. Configure the integration request mapping template with Content-Type of application/json and statusCode of 200. Configure the integration response mapping template with Content-Type of text/html. In the integration response mapping template, include the LandingPage HTML code that references the APIs.
Answer: (SHOW ANSWER)

NEW QUESTION: 36
A developer needs to retrieve all data from an Amazon DynamoDB table that matches a particular partition key.
Which solutions will meet this requirement in the MOST operationally efficient way? (Select TWO.)
A. Use the Scan API and a filter expression to match on the key.
B. Use the GetItem API and a PartiQL statement to match on the key.
C. Use the ExecuteStatement API and a filter expression to match on the key.
D. Use the ExecuteStatement API and a PartiQL statement to match on the key.
E. Use the GetItem API with a request parameter for key that contains the partition key name and specific key value.
Answer: D,E (LEAVE A REPLY)

NEW QUESTION: 37

A developer is building an application on AWS. The application has an Amazon API Gateway API that sends requests to an AWS Lambda function. The API is experiencing increased latency because the Lambda function has limited available CPU to fulfill the requests. Before the developer deploys the API into production, the developer must configure the Lambda function to have more CPU.
Which solution will meet this requirement?
A. Increase the virtual CPU (vCPU) cores quota of the Lambda function.
B. Increase the ephemeral storage size of the Lambda function.
C. Increase the amount of memory that is allocated to the Lambda function.
D. Increase the timeout value of the Lambda function.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 38
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: A,C (LEAVE A REPLY)
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.

NEW QUESTION: 39
A company is building a serverless application that uses AWS Lambda functions. The company needs to create a set of test events to test Lambda functions in a development environment. The test events will be created once and then will be used by all the developers in an 1AM developer group. The test events must be editable by any of the 1AM users in the 1AM developer group.

Which solution will meet these requirements?
A. Create and store the test events in Amazon DynamoDB. Allow access to DynamoDB by using 1AM roles.
B. Create the test events. Configure the event sharing settings to make the test events private.
C. Create the test events. Configure the event sharing settings to make the test events shareable.
D. Create and store the test events in Amazon S3 as JSON objects. Allow S3 bucket access to all 1AM users.
Answer: (SHOW ANSWER)

NEW QUESTION: 40
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 (LEAVE A REPLY)
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.

NEW QUESTION: 41
A company hosts a batch processing application on AWS Elastic Beanstalk with instances that run the most recent version of Amazon Linux. The application sorts and processes large datasets. In recent weeks, the application's performance has decreased significantly during a peak period for traffic. A developer suspects that the application issues are related to the memory usage. The developer checks the Elastic Beanstalk console and notices that memory usage is not being tracked.
How should the developer gather more information about the application performance issues?
A. Configure the Amazon CloudWatch agent to push logs to Amazon CloudWatch Logs by using port 443.

B. Configure the Elastic Beanstalk .ebextensions directory to track the memory usage of the instances.
C. Configure the Amazon CloudWatch agent to track the memory usage of the instances.
D. Configure an Amazon CloudWatch dashboard to track the memory usage of the instances.
Answer: C (LEAVE A REPLY)
Comprehensive Detailed Explanation with all AWS Reference
To monitor memory usage in Amazon Elastic Beanstalk environments, it's important to understand that default Elastic Beanstalk monitoring capabilities in Amazon CloudWatch do not track memory usage, as memory metrics are not collected by default. Instead, the Amazon CloudWatch agent must be configured to collect memory usage metrics.
Why Option C is Correct:
The Amazon CloudWatch agent can be installed and configured to monitor system-level metrics such as memory and disk utilization.
To enable memory tracking, developers need to install the CloudWatch agent on the Amazon Elastic Compute Cloud (EC2) instances associated with the Elastic Beanstalk environment. After installation, the agent can be configured to collect memory metrics, which can then be sent to CloudWatch for further analysis.
How to Implement This Solution:
Install the CloudWatch Agent:
Use .ebextensions or AWS Systems Manager to install and configure the CloudWatch agent on the EC2 instances running in the Elastic Beanstalk environment.
Modify CloudWatch Agent Configuration:
Create a config.json file that specifies memory usage tracking and other desired metrics.
Enable Metrics Reporting:
The CloudWatch agent can push the metrics to CloudWatch, where they can be monitored.
Why Other Options are Incorrect:
Option A: Configuring the agent to push logs is not sufficient to track memory metrics. This option addresses logging but not system-level metrics like memory usage.
Option B: The .ebextensions directory is used to customize Elastic Beanstalk environments but does not directly track memory metrics without additional configuration of the CloudWatch agent.
Option D: Configuring a CloudWatch dashboard will only visualize the metrics that are already being collected. It will not enable memory usage tracking.
AWS Documentation Reference:
Amazon CloudWatch Agent Overview
Elastic Beanstalk Customization Using .ebextensions
Monitoring Custom Metrics

NEW QUESTION: 42
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 policy. Store the access key ID and secret access key on the EC2 instances
C. Modify the application to use the S3 GeneratePresignedUrl API call
D. Modify the application to use the S3 GetObject API call and to return the object handle to the user
E. Modify the application to delegate requests to the S3 bucket.
Answer: A,C (LEAVE A REPLY)
IAM Roles for EC2 (A): The most secure way to provide AWS permissions from EC2.
Create a role with a policy allowing s3:GetObject on the specific bucket.
Attach the role to an instance profile and associate that profile with your instances.
Pre-signed URLs (C): Temporary, authenticated URLs for specific S3 actions.
Modify the app to use the AWS SDK to call GeneratePresignedUrl.
Embed these URLs when a user is properly logged in, allowing download access.
Reference:
IAM Roles for EC2: https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html Generating Presigned URLs:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/ShareObjectPreSignedURL.htm

NEW QUESTION: 43
A developer at a company recently created a serverless application to process and show data from business reports. The application's user interface (UI) allows users to select and start processing the files. The UI displays a message when the result is available to view. The application uses AWS Step Functions with AWS Lambda functions to process the files. The developer used Amazon API Gateway and Lambda functions to create an API to support the UI.
The company's UI team reports that the request to process a file is often returning timeout errors because of the see or complexity of the files. The UI team wants the API to provide an immediate response so that the UI can deploy a message while the files are being processed. The backend process that is invoked by the API needs to send an email message when the report processing is complete.
What should the developer do to configure the API to meet these requirements?
A. Change the API Gateway route to add an X-Amz-Invocation-Type header win a sialic value of 'Event' in the integration request Deploy the API Gateway stage to apply the changes.

B. Change the configuration of the Lambda function that implements the request to process a file. Configure the maximum age of the event so that the Lambda function will ion asynchronously.
C. Change the API Gateway timeout value to match the Lambda function ominous value. Deploy the API Gateway stage to apply the changes.
D. Change the API Gateway route to add an X-Amz-Target header with a static value of 'A sync' in the integration request Deploy me API Gateway stage to apply the changes.
Answer: A (LEAVE A REPLY)
This solution allows the API to invoke the Lambda function asynchronously, which means that the API will return an immediate response without waiting for the function to complete. The X-Amz-Invocation-Type header specifies the invocation type of the Lambda function, and setting it to 'Event' means that the function will be invoked asynchronously. The function can then use Amazon Simple Email Service (SES) to send an email message when the report processing is complete.

NEW QUESTION: 44
A company needs to harden its container images before the images are in a running state. The company's application uses Amazon Elastic Container Registry (Amazon ECR) as an image registry. Amazon Elastic Kubernetes Service (Amazon EKS) for compute, and an AWS CodePipeline pipeline that orchestrates a continuous integration and continuous delivery (CI/CD) workflow.
Dynamic application security testing occurs in the final stage of the pipeline after a new image is deployed to a development namespace in the EKS cluster. A developer needs to place an analysis stage before this deployment to analyze the container image earlier in the CI/CD pipeline.
Which solution will meet these requirements with the MOST operational efficiency?
A. Build the container image and run the docker scan command locally. Mitigate any findings before pushing changes to the source code repository. Write a pre-commit hook that enforces the use of this workflow before commit.
B. Create a new CodePipeline stage that occurs after the container image is built. Configure ECR basic image scanning to scan on image push. Use an AWS Lambda function as the action provider. Configure the Lambda function to check the scan results and to fail the pipeline if there are findings.
C. Create a new CodePipeline stage that occurs after source code has been retrieved from its repository. Run a security scanner on the latest revision of the source code. Fail the pipeline if there are findings.
D. Add an action to the deployment stage of the pipeline so that the action occurs before the deployment to the EKS cluster. Configure ECR basic image scanning to scan on image push. Use an AWS Lambda function as the action provider. Configure the Lambda function to check the scan results and to fail the pipeline if there are findings.
Answer: B (LEAVE A REPLY)

The solution that will meet the requirements with the most operational efficiency is to create a new CodePipeline stage that occurs after the container image is built. Configure ECR basic image scanning to scan on image push. Use an AWS Lambda function as the action provider. Configure the Lambda function to check the scan results and to fail the pipeline if there are findings. This way, the container image is analyzed earlier in the CI/CD pipeline and any vulnerabilities are detected and reported before deploying to the EKS cluster. The other options either delay the analysis until after deployment, which increases the risk of exposing insecure images, or perform analysis on the source code instead of the container image, which may not capture all the dependencies and configurations that affect the security posture of the image.

NEW QUESTION: 45
A company uses a custom root certificate authority certificate chain (Root CA Cert) that is 10 KB in size generate SSL certificates for its on-premises HTTPS endpoints. One of the company's cloud based applications has hundreds of AWS Lambda functions that pull date from these endpoints. A developer updated the trust store of the Lambda execution environment to use the Root CA Cert when the Lambda execution environment is initialized. The developer bundled the Root CA Cert as a text file in the Lambdas deployment bundle. After 3 months of development the root CA Cert is no longer valid and must be updated. The developer needs a more efficient solution to update the Root CA Cert for all deployed Lambda functions. The solution must not include rebuilding or updating all Lambda functions that use the Root CA Cert. The solution must also work for all development, testing and production environment. Each environment is managed in a separate AWS account.
When combination of steps Would the developer take to meet these environments MOST cost-effectively? (Select TWO)
A. Store the Root CA Cert as a secret in AWS Secrets Manager. Create a resource-based policy. Add IAM users to allow access to the secret
B. Store the Root CA Cert as a Secure Sting parameter in aws Systems Manager Parameter Store Create a resource-based policy. Add IAM users to allow access to the policy.
C. Store the Root CA Cert in an Amazon S3 bucket. Create a resource- based policy to allow access to the bucket.
D. Refactor the Lambda code to load the Root CA Cert from the Root CA Certs location. Modify the runtime trust store inside the Lambda function handler.
E. Refactor the Lambda code to load the Root CA Cert from the Root CA Cert's location. Modify the runtime trust store outside the Lambda function handler.
Answer: (SHOW ANSWER)
This solution will meet the requirements by storing the Root CA Cert as a Secure String parameter in AWS Systems Manager Parameter Store, which is a secure and scalable service for storing and managing configuration data and secrets. The resource-based policy will allow IAM users in different AWS accounts and environments to access the parameter without requiring cross-account roles or permissions. The Lambda code will be refactored to load the

Root CA Cert from the parameter store and modify the runtime trust store outside the Lambda function handler, which will improve performance and reduce latency by avoiding repeated calls to Parameter Store and trust store modifications for each invocation of the Lambda function. Option A is not optimal because it will use AWS Secrets Manager instead of AWS Systems Manager Parameter Store, which will incur additional costs and complexity for storing and managing a non-secret configuration data such as Root CA Cert. Option C is not optimal because it will deactivate the application secrets and monitor the application error logs temporarily, which will cause application downtime and potential data loss. Option D is not optimal because it will modify the runtime trust store inside the Lambda function handler, which will degrade performance and increase latency by repeating unnecessary operations for each invocation of the Lambda function.

NEW QUESTION: 46
A company has deployed an application on AWS Elastic Beanstalk. The company has configured the Auto Scaling group that is associated with the Elastic Beanstalk environment to have five Amazon EC2 instances. If the capacity is fewer than four EC2 instances during the deployment, application performance degrades. The company is using the all-at-once deployment policy.
What is the MOST cost-effective way to solve the deployment issue?
A. Change the Auto Scaling group to six desired instances.
B. Change the deployment policy to traffic splitting. Specify an evaluation time of 1 hour.
C. Change the deployment policy to rolling with additional batch. Specify a batch size of 1.
D. Change the deployment policy to rolling. Specify a batch size of 2.
Answer: (SHOW ANSWER)
This solution will solve the deployment issue by deploying the new version of the application to one new EC2 instance at a time, while keeping the old version running on the existing instances. This way, there will always be at least four instances serving traffic during the deployment, and no downtime or performance degradation will occur. Option A is not optimal because it will increase the cost of running the Elastic Beanstalk environment without solving the deployment issue. Option B is not optimal because it will split the traffic between two versions of the application, which may cause inconsistency and confusion for the customers. Option D is not optimal because it will deploy the new version of the application to two existing instances at a time, which may reduce the capacity below four instances during the deployment.

Valid DVA-C02 Dumps shared by TrainingQuiz.com for Helping Passing DVA-C02 Exam! TrainingQuiz.com now offer the **newest DVA-C02 exam dumps**, the TrainingQuiz.com DVA-C02 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C02 dumps with Test Engine here:

https://www.trainingquiz.com/DVA-C02-practice-quiz.html (**649** Q&As Dumps, **40%OFF**
Special Discount: Exam-Tests)

NEW QUESTION: 47
A company is planning to deploy an application on AWS behind an Elastic Load Balancing (ELB) load balancer. The application uses an HTTP/HTTPS listener and must access the client IP addresses.
Which load-balancing solution meets these requirements?
A. Use a Network Load Balancer (NLB). Enable proxy protocol support on the NLB and the target application.
B. Use an Application Load Balancer and the X-Forwarded-For headers.
C. Use an Application Load Balancer. Register the targets by the instance ID.
D. Use a Network Load Balancer and the X-Forwarded-For headers.
Answer: (SHOW ANSWER)

NEW QUESTION: 48
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 Manager. Set the user name password, database, host and port. Turn on secret rotation. Create encrypted Lambda environment variables for the DynamoDB table, S3 bucket and SNS topic.
B. Create an RDS database secret in AWS Secrets Manager. Set the user name password, database, host and port. Turn on secret rotation. Create Secure String parameters in AWS Systems Manager Parameter Store for the DynamoDB table, S3 bucket and SNS topic.
C. Create RDS database parameters in AWS Systems Manager Parameter. Store for the user name password, database, host and port. Create encrypted Lambda environment variables for me DynamoDB table, S3 bucket, and SNS topic. Create a Lambda function and set the logic for the credentials rotation task Schedule the credentials rotation task in Amazon EventBridge.
D. Create RDS database parameters in AWS Systems Manager Parameter. Store for the user name password database, host, and port. Store the DynamoDB table. 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 (LEAVE A REPLY)

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.

NEW QUESTION: 49
An online food company provides an Amazon API Gateway HTTP API 1o receive orders for partners. The API is integrated with an AWS Lambda function. The Lambda function stores the orders in an Amazon DynamoDB table.
The company expects to onboard additional partners Some to me panthers require additional Lambda function to receive orders. The company has created an Amazon S3 bucket. The company needs 10 store all orders and updates m the S3 bucket for future analysis How can the developer ensure that an orders and updates are stored to Amazon S3 with the LEAST development effort?
A. Create a new Lambda function and a new API Gateway API endpoint. Configure the new Lambda function to write to the S3 bucket. Modify the original Lambda function to post updates to the new API endpoint.
B. Use Amazon Kinesis Data Streams to create a new data stream. Modify the Lambda function to publish orders to the oats stream Configure in data stream to write to the S3 bucket.
C. Enable DynamoDB Streams on me DynamoOB table. Create a new lambda function. Associate the stream's Amazon Resource Name (ARN) with the Lambda Function Configure the Lambda function to write to the S3 bucket as records appear in the table's stream.
D. Modify the Lambda function to punish to a new Amazon. Simple Lambda function receives orders. Subscribe a new Lambda function to the topic. Configure the new Lambda function to write to the S3 bucket as updates come through the topic.
Answer: (SHOW ANSWER)

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

NEW QUESTION: 50
A company is building a content authoring application. The application has multiple user groups, such as content creator, reviewer, approver, and administrator. The company needs to assign users fine-grained permissions for specific parts of the application.
The company needs a solution to configure, maintain, and analyze user permissions. The company wants a solution that can be easily adapted to work with newer applications in the future. The company must use a third-party OpenID Connect (OIDC) identity provider (IdP) to authenticate users.
A. Configure an Amazon Cognito identity pool for the application. Use the identity pool identities within the application to manage user permissions.
B. Configure the application to check user permissions upon request. Configure the application logic to manage user permissions.
C. Use Amazon Verified Permissions to set up user permissions. Integrate Verified Permissions with a third-party IdP. Configure the application to request authorization decisions from Verified Permissions.
D. Set up an IAM role for each user group. Assign users appropriate IAM roles. Configure the application to determine appropriate permissions for each user based on the user's IAM role.
Answer: C (LEAVE A REPLY)
Comprehensive Detailed Explanation with all AWS Reference
Why Option C is Correct:
Amazon Verified Permissions provides fine-grained access control capabilities, making it ideal for managing complex user permissions. It integrates with OIDC IdPs for authentication and allows applications to request authorization decisions dynamically. It is also easily adaptable to newer applications.
Why Other Options are Incorrect:
Option A: Cognito identity pools do not natively support fine-grained permission analysis or management.
Option B: Managing permissions in application logic adds significant operational overhead.

Option D: IAM roles are not designed for application-specific fine-grained access control and are more suitable for resource-level permissions.
AWS Documentation Reference:
Amazon Verified Permissions

NEW QUESTION: 51
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: (SHOW ANSWER)
Distributed Tracing with X-Ray: X-Ray helps visualize request paths and identify bottlenecks in applications distributed across Regions.
Region Annotations (Automatic for AWS Services): X-Ray automatically adds a Region annotation to segments representing calls to AWS services. This aids in tracing cross-Region traffic.
Region Annotations (Manual for User-Defined): For segments representing calls to user-defined services in different Regions, the developer needs to add the Region annotation manually to enable comprehensive tracing.
Reference:
AWS X-Ray: https://aws.amazon.com/xray/

NEW QUESTION: 52
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 accounts. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the API. Use the Lambda function to store

the photos and details in the DynamoDB table. Retrieve previously uploaded photos directly from the DynamoDB table.
B. Use Amazon Cognito user pools to manage user accounts. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the API. 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 table. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
C. Create an IAM user for each user of the application during the sign-up process. Use IAM authentication to access the API Gateway API. 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 table. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
D. Create a users table in DynamoDB. Use the table to manage user accounts. Create a Lambda authorizer that validates user credentials against the users table. Integrate the Lambda authorizer with API Gateway to control access to the API. 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 table. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
Answer: B (LEAVE A REPLY)
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.
Reference:
[Amazon Cognito User Pools]
[Use Amazon Cognito User Pools - Amazon API Gateway]
[Amazon Simple Storage Service (S3)]
[Amazon DynamoDB]

NEW QUESTION: 53
A bookstore has an ecommerce website that stores order information in an Amazon DynamoDB table named BookOrders. The DynamoDB table contains approximately one million records.
The table uses OrderID as a partition key. There are no other indexes.
A developer wants to build a new reporting feature to retrieve all records from the table for a specified customer, based on a CustomerID property.
A. Create a DynamoDB global secondary index (GSI) on the table. Use CustomerID as the partition key. Use the specified CustomerID value to run a query on the table.

B. Create a DynamoDB global secondary index (GSI) on the table. Use CustomerID as the sort key. Use a filter expression to perform a scan operation on the table to match on the specified CustomerID value.
C. Create a DynamoDB local secondary index (LSI) on the table. Use CustomerID as the sort key. Run a PartiQL query on the table with a SELECT statement where CustomerID equals the specified CustomerID value.
D. Create a DynamoDB local secondary index (LSI) on the table. Use CustomerID as the partition key. Use the specified CustomerID value to run a query on the table.
Answer: (SHOW ANSWER)
Comprehensive and Detailed Step-by-Step
The requirement is to query records by CustomerID, which is not the current partition key (OrderID). To achieve this efficiently:
Option A: Create a GSI with CustomerID as the Partition Key:
A Global Secondary Index (GSI) allows developers to create a different partition key and optional sort key for querying the data.
By creating a GSI with CustomerID as the partition key, the developer can query the table efficiently using CustomerID as the primary lookup key.
This avoids scanning the entire table and matches the requirement.
Why Other Options Are Incorrect:
Option B: Using CustomerID as a sort key for the GSI and performing a scan operation is inefficient. Queries are optimized, but scans are not.
Option C and D: Local Secondary Indexes (LSI) are only valid when the partition key remains the same as the base table. Since OrderID is the base table's partition key, using CustomerID as the partition key or sort key in an LSI is not valid.
Reference:
Amazon DynamoDB Documentation: GSIs

NEW QUESTION: 54
An AWS Lambda function requires read access to an Amazon S3 bucket and requires read/write access to an Amazon DynamoDB table The correct 1AM policy already exists What is the MOST secure way to grant the Lambda function access to the S3 bucket and the DynamoDB table?
A. Attach the existing 1AM policy to the Lambda function.
B. Create an 1AM role for the Lambda function Attach the existing 1AM policy to the role Attach the role to the Lambda function
C. Create an 1AM user with programmatic access Attach the existing 1AM policy to the user. Add the user access key ID and secret access key as environment variables in the Lambda function.
D. Add the AWS account root user access key ID and secret access key as encrypted environment variables in the Lambda function
Answer: B (LEAVE A REPLY)

Principle of Least Privilege: Granting specific permissions through an IAM role is more secure than directly attaching policies to a function or using root user credentials.
IAM Roles for Lambda: Designed to provide temporary credentials to Lambda functions, enhancing security.
Reusability: The existing IAM policy ensures the correct S3 and DynamoDB access is granted.
Reference:
IAM Roles for Lambda Documentation:
https://docs.aws.amazon.com/lambda/latest/dg/lambda-intro-execution-role.html IAM Best Practices: https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html

NEW QUESTION: 55
A developer created an AWS Lambda function that performs a series of operations that involve multiple AWS services. The function's duration time is higher than normal. To determine the cause of the issue, the developer must investigate traffic between the services without changing the function code Which solution will meet these requirements?
A. Enable AWS X-Ray active tracing in the Lambda function Review the logs in X-Ray
B. Configure AWS CloudTrail View the trail logs that are associated with the Lambda function.
C. Review the AWS Config logs in Amazon Cloud Watch.
D. Review the Amazon CloudWatch logs that are associated with the Lambda function.
Answer: (SHOW ANSWER)
Tracing Distributed Systems: AWS X-Ray is designed to trace requests across services, helping identify bottlenecks in distributed applications like this one.
No Code Changes: Enabling X-Ray tracing often requires minimal code changes, meeting the requirement.
Identifying Bottlenecks: Analyzing X-Ray traces and logs will reveal latency in communications between different AWS services, leading to the high duration time.
Reference:
AWS X-Ray: https://aws.amazon.com/xray/
X-Ray and Lambda: https://docs.aws.amazon.com/xray/latest/devguide/xray-services-lambda.html

NEW QUESTION: 56
A company has an application that runs as a series of AWS Lambda functions. Each Lambda function receives data from an Amazon Simple Notification Service (Amazon SNS) topic and writes the data to an Amazon Aurora DB instance.
To comply with an information security policy, the company must ensure that the Lambda functions all use a single securely encrypted database connection string to access Aurora.
Which solution will meet these requirements'?
A. Use IAM database authentication for Aurora to enable secure database connections for ail the Lambda functions.
B. Store the credentials and read the credentials from an encrypted Amazon RDS DB instance.

C. Store the credentials in AWS Systems Manager Parameter Store as a secure string parameter.
D. Use Lambda environment variables with a shared AWS Key Management Service (AWS KMS) key for encryption.
Answer: A (LEAVE A REPLY)
This solution will meet the requirements by using IAM database authentication for Aurora, which enables using IAM roles or users to authenticate with Aurora databases instead of using passwords or other secrets. The developer can use IAM database authentication for Aurora to enable secure database connections for all the Lambda functions that access Aurora DB instance. The developer can create an IAM role with permission to connect to Aurora DB instance and attach it to each Lambda function. The developer can also configure Aurora DB instance to use IAM database authentication and enable encryption in transit using SSL certificates. This way, the Lambda functions can use a single securely encrypted database connection string to access Aurora without needing any secrets or passwords. Option B is not optimal because it will store the credentials and read them from an encrypted Amazon RDS DB instance, which may introduce additional costs and complexity for managing and accessing another RDS DB instance. Option C is not optimal because it will store the credentials in AWS Systems Manager Parameter Store as a secure string parameter, which may require additional steps or permissions to retrieve and decrypt the credentials from Parameter Store. Option D is not optimal because it will use Lambda environment variables with a shared AWS Key Management Service (AWS KMS) key for encryption, which may not be secure or scalable as environment variables are stored as plain text unless encrypted with AWS KMS.

NEW QUESTION: 57
A developer deployed an application to an Amazon EC2 instance The application needs to know the public IPv4 address of the instance How can the application find this information?
A. Query the instance metadata from http./M69.254.169.254. latestmeta-data/.
B. Query the instance user data from http '169 254.169 254. latest/user-data/
C. Query the Amazon Machine Image (AMI) information from http://169.254.169.254/latest/meta-data/ami/.
D. Check the hosts file of the operating system
Answer: A (LEAVE A REPLY)
Instance Metadata Service: EC2 instances have access to an internal metadata service. It provides instance-specific information like instance ID, security groups, and public IP address.
Accessing Metadata:
Make an HTTP GET request to the base URL: http://169.254.169.254/latest/meta-data/ You'll get a list of available categories. The public IPv4 address is under public-ipv4.
Reference:
Instance Metadata and User Data:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-retrieval.html

NEW QUESTION: 58
A company is working on a new serverless application. A developer needs to find an automated way to deploy AWS Lambda functions and the dependent Infrastructure with minimum coding effort. The application also needs to be reliable.
Which method will meet these requirements with the LEAST operational overhead?
A. Build the application by using shell scripts to create .zip files for each Lambda function. Manually upload the .zip files to the AWS Management Console.
B. Build the application by using the AWS Serverless Application Model (AWS SAM). Use a continuous integration and continuous delivery (CI/CD) pipeline and the SAM CLI to deploy the Lambda functions.
C. Build a container for each Lambda function. Store the container images in AWS CodeArtifact. Deploy the containers as Lambda functions by using the AWS CLI in a continuous integration and continuous delivery (CI/CD) pipeline.
D. Build the application by using shell scripts to create .zip files for each Lambda function. Upload the .zip files. Deploy the .zip files as Lambda functions by using the AWS CLI in a continuous integration and continuous delivery (CI/CD) pipeline.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 59
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 emad_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 emal_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 (LEAVE A REPLY)

Understand the Problem: The existing DynamoDB table has email_address as the partition key. Searching by customer_type requires a different data access pattern. We need an efficient way to query for partial matches on email_address based on customer_type.
Why Global Secondary Index (GSI):
GSIs allow you to define a different partition key and sort key from the main table, enabling new query patterns.
In this case, having customer_type as the GSI's partition key lets you group all emails with the same customer type together.
Using email_address as the sort key allows ordering within each customer type, facilitating the partial matching.
Querying the GSI:
You'll perform a query operation on the GSI, not the original table.
Use the begins_with key condition expression on the GSI's sort key (email_address) to find partial matches as the user types in the customer_type field.
Reference:
DynamoDB Global Secondary Indexes:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GSI.html DynamoDB Query Operation: [invalid URL removed] Key Condition Expressions: [invalid URL removed]

NEW QUESTION: 60
A company stores customer credit reports in an Amazon S3 bucket. An analytics service uses standard Amazon S3 GET requests to access the reports. A developer must implement a solution to redact personally identifiable information (PII) from the reports before the reports reach the analytics service.
A. Load the S3 objects into Amazon Redshift by using a COPY command. Implement dynamic data masking. Refactor the analytics service to read from Amazon Redshift.
B. Set up an S3 Object Lambda function. Attach the function to an S3 Object Lambda Access Point. Program the function to call a PII redaction API.
C. Use AWS Key Management Service (AWS KMS) to implement encryption in the S3 bucket. Re-upload all the existing S3 objects. Give the kms permission to the analytics service.
D. Create an Amazon Simple Notification Service (Amazon SNS) topic. Implement message data protection. Refactor the analytics service to publish data access requests to the SNS topic.
Answer: B (LEAVE A REPLY)
Comprehensive Detailed Step by Step Explanation with All AWS Developer Reference:
To redact PII from S3 objects before they are accessed by the analytics service, the most efficient solution is to use S3 Object Lambda. S3 Object Lambda allows you to add your own code (Lambda function) to process and transform data when it is retrieved from Amazon S3. You can attach a Lambda function to an S3 Object Lambda Access Point, which in this case would run a redaction API to remove PII from the reports.

Operational Efficiency: S3 Object Lambda handles data processing on the fly, without requiring the data to be permanently transformed or moved to another service (like Amazon Redshift).
Alternatives:
Option A: Loading the data into Amazon Redshift would require refactoring the analytics service and maintaining an additional data pipeline, increasing complexity.
Option C: Using AWS KMS for encryption protects data at rest and in transit, but it does not address PII redaction.
Option D: SNS is a messaging service and does not support direct data transformation.

NEW QUESTION: 61
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 application. Deploy the unit tests by using the amplify push command in the Amplify CLI.
C. Add a test phase to the amplify.yml build settings for the application.
D. Add a test phase to the aws-exports.js file for the application.
Answer: (SHOW ANSWER)
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.

Valid DVA-C02 Dumps shared by TrainingQuiz.com for Helping Passing DVA-C02 Exam! TrainingQuiz.com now offer the **newest DVA-C02 exam dumps**, the TrainingQuiz.com DVA-C02 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C02 dumps with Test Engine here: https://www.trainingquiz.com/DVA-C02-practice-quiz.html (**649** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

NEW QUESTION: 62
A developer has an application that makes batch requests directly to Amazon DynamoDB by using the BatchGetItem low-level API operation. The responses frequently return values in the UnprocessedKeys element.
Which actions should the developer take to increase the resiliency of the application when the batch response includes values in UnprocessedKeys? (Choose two.)

A. Retry the batch operation immediately.
B. Retry the batch operation with exponential backoff and randomized delay.
C. Update the application to use an AWS software development kit (AWS SDK) to make the requests.
D. Increase the provisioned read capacity of the DynamoDB tables that the operation accesses.
E. Increase the provisioned write capacity of the DynamoDB tables that the operation accesses.
Answer: B,C (LEAVE A REPLY)
The UnprocessedKeys element indicates that the BatchGetItem operation did not process all of the requested items in the current response. This can happen if the response size limit is exceeded or if the table's provisioned throughput is exceeded. To handle this situation, the developer should retry the batch operation with exponential backoff and randomized delay to avoid throttling errors and reduce the load on the table. The developer should also use an AWS SDK to make the requests, as the SDKs automatically retry requests that return UnprocessedKeys.
Reference:
[BatchGetItem - Amazon DynamoDB]
[Working with Queries and Scans - Amazon DynamoDB]
[Best Practices for Handling DynamoDB Throttling Errors]

NEW QUESTION: 63
An application that runs on AWS Lambda requires access to specific highly confidential objects in an Amazon S3 bucket. In accordance with the principle of least privilege a company grants access to the S3 bucket by using only temporary credentials.
How can a developer configure access to the S3 bucket in the MOST secure way?
A. Hardcode the credentials that are required to access the S3 objects in the application code. Use the credentials to access me required S3 objects.
B. Create a secret access key and access key ID with permission to access the S3 bucket. Store the key and key ID in AWS Secrets Manager. Configure the application to retrieve the Secrets Manager secret and use the credentials to access me S3 objects.
C. Create a Lambda function execution role Attach a policy to the rote that grants access to specific objects in the S3 bucket.
D. Create a secret access key and access key ID with permission to access the S3 bucket Store the key and key ID as environment variables m Lambda. Use the environment variables to access the required S3 objects.
Answer: C (LEAVE A REPLY)
This solution will meet the requirements by creating a Lambda function execution role, which is an IAM role that grants permissions to a Lambda function to access AWS resources such as Amazon S3 objects. The developer can attach a policy to the role that grants access to specific objects in the S3 bucket that are required by the application, following the principle of least

privilege. Option A is not optimal because it will hardcode the credentials that are required to access S3 objects in the application code, which is insecure and difficult to maintain. Option B is not optimal because it will create a secret access key and access key ID with permission to access the S3 bucket, which will introduce additional security risks and complexity for storing and managing credentials. Option D is not optimal because it will store the secret access key and access key ID as environment variables in Lambda, which is also insecure and difficult to maintain.

NEW QUESTION: 64
A company is implementing an application on Amazon EC2 instances. The application needs to process incoming transactions. When the application detects a transaction that is not valid, the application must send a chat message to the company's support team. To send the message, the application needs to retrieve the access token to authenticate by using the chat API.
A developer needs to implement a solution to store the access token. The access token must be encrypted at rest and in transit. The access token must also be accessible from other AWS accounts.
Which solution will meet these requirements with the LEAST management overhead?
A. Use an AWS Systems Manager Parameter Store SecureString parameter that uses an AWS Key Management Service (AWS KMS) AWS managed key to store the access token. Add a resource-based policy to the parameter to allow access from other accounts. Update the IAM role of the EC2 instances with permissions to access Parameter Store. Retrieve the token from Parameter Store with the decrypt flag enabled. Use the decrypted access token to send the message to the chat.
B. Encrypt the access token by using an AWS Key Management Service (AWS KMS) customer managed key. Store the access token in an Amazon DynamoDB table. Update the IAM role of the EC2 instances with permissions to access DynamoDB and AWS KMS. Retrieve the token from DynamoDB. Decrypt the token by using AWS KMS on the EC2 instances. Use the decrypted access token to send the message to the chat.
C. Use AWS Secrets Manager with an AWS Key Management Service (AWS KMS) customer managed key to store the access token. Add a resource-based policy to the secret to allow access from other accounts. Update the IAM role of the EC2 instances with permissions to access Secrets Manager. Retrieve the token from Secrets Manager. Use the decrypted access token to send the message to the chat.
D. Encrypt the access token by using an AWS Key Management Service (AWS KMS) AWS managed key. Store the access token in an Amazon S3 bucket. Add a bucket policy to the S3 bucket to allow access from other accounts. Update the IAM role of the EC2 instances with permissions to access Amazon S3 and AWS KMS. Retrieve the token from the S3 bucket. Decrypt the token by using AWS KMS on the EC2 instances. Use the decrypted access token to send the massage to the chat.
Answer: C (LEAVE A REPLY)

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

NEW QUESTION: 65
A developer is preparing to begin development of a new version of an application. The previous version of the application is deployed in a production environment. The developer needs to deploy fixes and updates to the current version during the development of the new version of the application. The code for the new version of the application is stored in AWS CodeCommit.
Which solution will meet these requirements?
A. From the main branch, create a feature branch for production bug fixes. Create a second feature branch from the main branch for development of the new version.
B. Create a Git tag of the code that is currently deployed in production. Create a Git tag for the development of the new version. Push the two tags to the CodeCommit repository.
C. From the main branch, create a branch of the code that is currently deployed in production. Apply an IAM policy that ensures no other other users can push or merge to the branch.
D. Create a new CodeCommit repository for development of the new version of the application. Create a Git tag for the development of the new version.
Answer: (SHOW ANSWER)
A feature branch is a branch that is created from the main branch to work on a specific feature or task1. Feature branches allow developers to isolate their work from the main branch and avoid conflicts with other changes1. Feature branches can be merged back to the main branch when the feature or task is completed and tested1.
In this scenario, the developer needs to maintain two parallel streams of work: one for fixing and updating the current version of the application that is deployed in production, and another for developing the new version of the application. The developer can use feature branches to achieve this goal.
The developer can create a feature branch from the main branch for production bug fixes. This branch will contain the code that is currently deployed in production, and any fixes or updates that need to be applied to it. The developer can push this branch to the CodeCommit repository and use it to deploy changes to the production environment.
The developer can also create a second feature branch from the main branch for development of the new version of the application. This branch will contain the code that is under development for the new version, and any changes or enhancements that are part of it. The developer can push this branch to the CodeCommit repository and use it to test and deploy the new version of the application in a separate environment.
By using feature branches, the developer can keep the main branch stable and clean, and avoid mixing code from different versions of the application. The developer can also easily switch between branches and merge them when needed.

NEW QUESTION: 66
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: (SHOW ANSWER)
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.
Reference:
[Amazon DynamoDB]
[Amazon Simple Storage Service (S3)]

NEW QUESTION: 67
An organization is using Amazon CloudFront to ensure that its users experience low-latency access to its web application. The organization has identified a need to encrypt all traffic between users and CloudFront, and all traffic between CloudFront and the web application.
How can these requirements be met? (Select TWO)
A. Use AWS KMS t0 encrypt traffic between cloudFront and the web application.
B. Set the Origin Protocol Policy to "HTTPS Only".
C. Set the Origin's HTTP Port to 443.
D. Set the Viewer Protocol Policy to "HTTPS Only" or Redirect HTTP to HTTPS"
E. Enable the CloudFront option Restrict Viewer Access.
Answer: (SHOW ANSWER)
This solution will meet the requirements by ensuring that all traffic between users and CloudFront, and all traffic between CloudFront and the web application, are encrypted using HTTPS protocol. The Origin Protocol Policy determines how CloudFront communicates with the origin server (the web application), and setting it to "HTTPS Only" will force CloudFront to

use HTTPS for every request to the origin server. The Viewer Protocol Policy determines how CloudFront responds to HTTP or HTTPS requests from users, and setting it to "HTTPS Only" or "Redirect HTTP to HTTPS" will force CloudFront to use HTTPS for every response to users. Option A is not optimal because it will use AWS KMS to encrypt traffic between CloudFront and the web application, which is not necessary or supported by CloudFront. Option C is not optimal because it will set the origin's HTTP port to 443, which is incorrect as port 443 is used for HTTPS protocol, not HTTP protocol. Option E is not optimal because it will enable the CloudFront option Restrict Viewer Access, which is used for controlling access to private content using signed URLs or signed cookies, not for encrypting traffic.

NEW QUESTION: 68
A developer is creating a stock trading application. The developer needs a solution to send text messages to application users to confirmation when a trade has been completed.
The solution must deliver messages in the order a user makes stock trades. The solution must not send duplicate messages.
Which solution will meet these requirements?
A. Configure a pipe in Amazon EventBridge Pipes. Connect the application to the pipe as a source. Configure the pipe to use each user's mobile phone number as a target. Configure the pipe to send incoming events to the users.
B. Configure the application to publish messages to an Amazon Data Firehose delivery stream. Configure the delivery stream to have a destination of each user's mobile phone number that is passed in the trade confirmation message.
C. Create an Amazon Simple Notification Service (SNS) FIFO topic. Configure the application to use the AWS SDK to publish notifications to the SNS topic to send SMS messages to the users.
D. Create an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Use the SendMessageIn API call to send the trade confirmation messages to the queue. Use the SendMessageOut API to send the messages to users by using the information provided in the trade confirmation message.
Answer: (SHOW ANSWER)

NEW QUESTION: 69
A company has a website that displays a daily newsletter. When a user visits the website, an AWS Lambda function processes the browser's request and queries the company's on-premises database to obtain the current newsletter. The newsletters are stored in English. The Lambda function uses the Amazon Translate TranslateText API operation to translate the newsletters, and the translation is displayed to the user.
Due to an increase in popularity, the website's response time has slowed. The database is overloaded. The company cannot change the database and needs a solution that improves the response time of the Lambda function.
Which solution meets these requirements?

A. Enable TranslateText API caching.
B. Cache the translated newsletters in the Lambda /tmp directory.
C. Change the Lambda function to use parallel processing.
D. Change to asynchronous Lambda function invocation.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 70
A developer is building an application that processes a stream of user-supplied data. The data stream must be consumed by multiple Amazon EC2 based processing applications in parallel and in real time. Each processor must be able to resume without losing data if there is a service interruption. The application architect plans to add other processors in the near future, and wants to minimize the amount of data duplication involved.
Which solution will satisfy these requirements?
A. Publish the data to Amazon Kinesis Data Streams.
B. Publish the data to Amazon Data Firehose.
C. Publish the data to Amazon Simple Queue Service (Amazon SQS).
D. Publish the data to Amazon EventBridge.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 71
A developer has code that is stored in an Amazon S3 bucket. The code must be deployed as an AWS Lambda function across multiple accounts in the same AWS Region as the S3 bucket an AWS CloudPormation template that runs for each account will deploy the Lambda function. What is the MOST secure way to allow CloudFormaton to access the Lambda Code in the S3 bucket?
A. Grant the CloudFormation service role the S3 ListBucket and GetObject permissions. Add a bucket policy to Amazon S3 with the principal of "AWS" (account numbers)
B. Grant the CloudFormation service row the S3 GetObfect permission. Add a Bucket policy to Amazon S3 with the principal of """
C. Use a service-based link to grant the Lambda function the S3 ListBucket and GetObject permissions by explicitly adding the S3 bucket's account number in the resource.
D. Use a service-based link to grant the Lambda function the S3 GetObject permission Add a resource of "** to allow access to the S3 bucket.
Answer: B (LEAVE A REPLY)
This solution allows the CloudFormation service role to access the S3 bucket from any account, as long as it has the S3 GetObject permission. The bucket policy grants access to any principal with the GetObject permission, which is the least privilege needed to deploy the Lambda code. This is more secure than granting ListBucket permission, which is not required for deploying Lambda code, or using a service-based link, which is not supported for Lambda functions.

NEW QUESTION: 72
A developer is building an application that uses an Amazon RDS for PostgreSQL database. To meet security requirements, the developer needs to ensure that data is encrypted at rest. The developer must be able to rotate the encryption keys on demand.
A. Use an AWS KMS managed encryption key to encrypt the database.
B. Create a symmetric customer managed AWS KMS key. Use the key to encrypt the database.
C. Create a 256-bit AES-GCM encryption key. Store the key in AWS Secrets Manager, and enable managed rotation. Use the key to encrypt the database.
D. Create a 256-bit AES-GCM encryption key. Store the key in AWS Secrets Manager. Configure an AWS Lambda function to perform key rotation. Use the key to encrypt the database.
Answer: B (LEAVE A REPLY)
Comprehensive Detailed Explanation with all AWS Reference
Why Option B is Correct:
A customer-managed AWS Key Management Service (KMS) key allows for encryption at rest and provides the ability to rotate the key on demand. This ensures compliance with security requirements for key management and database encryption.
RDS integrates natively with AWS KMS, allowing the use of a customer-managed key for encrypting data at rest.
Key rotation can be managed directly in AWS KMS without needing custom solutions.
Why Other Options are Incorrect:
Option A: AWS KMS managed encryption keys (AWS-owned keys) do not support key rotation on demand.
Option C & D: Storing keys in AWS Secrets Manager with custom rotation is not a recommended approach for database encryption. AWS KMS is designed specifically for secure key management and encryption.
AWS Documentation Reference:
Encrypting Amazon RDS Resources
AWS Key Management Service (KMS)

NEW QUESTION: 73
A developer has created a large AWS Lambda function. Deployment of the function Is failing because of an InvalidParameterValueException error. The error message indicates that the unzipped size of the function exceeds the maximum supported value.
Which actions can the developer take to resolve this error? (Select TWO.)
A. Use a compression algorithm that is more efficient than ZIP.
B. Break up the function into multiple smaller functions.
C. Zip the .zip file twice to compress the file more.
D. Move common libraries, function dependencies, and custom runtimes into Lambda layers.

E. Submit a quota increase request to AWS Support to increase the function to the required size.
Answer: (SHOW ANSWER)

NEW QUESTION: 74
A company caches session information for a web application in an Amazon DynamoDB table. The company wants an automated way to delete old items from the table.
What is the simplest way to do this?
A. Add an attribute with the expiration time; enable the Time To Live feature based on that attribute.
B. Add an attribute with the expiration time; name the attribute ItemExpiration.
C. Each day, create a new table to hold session data; delete the previous day's table.
D. Write a script that deletes old records; schedule the script as a cron job on an Amazon EC2 instance.
Answer: (SHOW ANSWER)

NEW QUESTION: 75
An application is using Amazon Cognito user pools and identity pools for secure access. A developer wants to integrate the user-specific file upload and download features in the application with Amazon S3. The developer must ensure that the files are saved and retrieved in a secure manner and that users can access only their own files. The file sizes range from 3 KB to 300 MB.
Which option will meet these requirements with the HIGHEST level of security?
A. Use S3 Event Notifications to validate the file upload and download requests and update the user interface (UI).
B. Save the details of the uploaded files in a separate Amazon DynamoDB table. Filter the list of files in the user interface (UI) by comparing the current user ID with the user ID associated with the file in the table.
C. Use Amazon API Gateway and an AWS Lambda function to upload and download files. Validate each request in the Lambda function before performing the requested operation.
D. Use an IAM policy within the Amazon Cognito identity prefix to restrict users to use their own folders in Amazon S3.
Answer: D (LEAVE A REPLY)
https://docs.aws.amazon.com/cognito/latest/developerguide/amazon-cognito-integrating-user-pools-with-identity-pools.html

NEW QUESTION: 76
A developer is managing an application that uploads user files to an Amazon S3 bucket named companybucket. The company wants to maintain copies of all the files uploaded by users for compliance purposes, while ensuring users still have access to the data through the application.

Which IAM permissions should be applied to users to ensure they can create but not remove files from the bucket?

A. json

Copy code

```
{
"Version": "2012-10-17",
"Statement": [
{
"Sid": "statement1",
"Effect": "Allow",
"Action": ["s3:GetObject", "s3:PutObject", "s3:DeleteObject"],
"Resource": ["arn:aws:s3:::companybucket"]
}
]
}
```

B. json

Copy code

```
{
"Version": "2012-10-17",
"Statement": [
{
"Sid": "statement1",
"Effect": "Allow",
"Action": ["s3:CreateBucket", "s3:GetBucketLocation"],
"Resource": "arn:aws:s3:::companybucket"
}
]
}
```

C. json

Copy code

```
{
"Version": "2012-10-17",
"Statement": [
{
"Sid": "statement1",
"Effect": "Allow",
"Action": ["s3:GetObject", "s3:PutObject", "s3:DeleteObject", "s3:PutObjectRetention"],
"Resource": "arn:aws:s3:::companybucket"
}
]
}
```

D. json

Copy code

```
{
"Version": "2012-10-17",
"Statement": [
{
"Sid": "statement1",
"Effect": "Allow",
"Action": ["s3:GetObject", "s3:PutObject"],
"Resource": ["arn:aws:s3:::companybucket"]
}
]
}
```

Answer: D (LEAVE A REPLY)

To meet the requirement:

Users must be able to upload (PutObject) and read (GetObject) files but not delete them.

Option D ensures users cannot delete files by omitting the s3:DeleteObject action while allowing s3:GetObject and s3:PutObject.

Option A: Includes s3:DeleteObject, which allows users to delete files and does not meet the requirement.

Option B: Contains unrelated actions like CreateBucket, which is not relevant here.

Option C: Adds s3:PutObjectRetention, which is unnecessary and does not restrict DeleteObject.

Reference:

AWS S3 Permissions Documentation

Reference:

AWS S3 Permissions Documentation

Valid DVA-C02 Dumps shared by TrainingQuiz.com for Helping Passing DVA-C02 Exam! TrainingQuiz.com now offer the **newest DVA-C02 exam dumps**, the TrainingQuiz.com DVA-C02 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C02 dumps with Test Engine here: https://www.trainingquiz.com/DVA-C02-practice-quiz.html (**649** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

NEW QUESTION: 77

A company is running Amazon EC2 instances in multiple AWS accounts. A developer needs to implement an application that collects all the lifecycle events of the EC2 instances. The

application needs to store the lifecycle events in a single Amazon Simple Queue Service (Amazon SQS) queue in the company's main AWS account for further processing.
Which solution will meet these requirements?
A. Configure Amazon EC2 to deliver the EC2 instance lifecycle events from all accounts to the Amazon EventBridge event bus of the main account. Add an EventBridge rule to the event bus of the main account that matches all EC2 instance lifecycle events. Add the SQS queue as a target of the rule.
B. Use the resource policies of the SQS queue in the main account to give each account permissions to write to that SQS queue. Add to the Amazon EventBridge event bus of each account an EventBridge rule that matches all EC2 instance lifecycle events. Add the SQS queue in the main account as a target of the rule.
C. Write an AWS Lambda function that scans through all EC2 instances in the company accounts to detect EC2 instance lifecycle changes. Configure the Lambda function to write a notification message to the SQS queue in the main account if the function detects an EC2 instance lifecycle change. Add an Amazon EventBridge scheduled rule that invokes the Lambda function every minute.
D. Configure the permissions on the main account event bus to receive events from all accounts. Create an Amazon EventBridge rule in each account to send all the EC2 instance lifecycle events to the main account event bus. Add an EventBridge rule to the main account event bus that matches all EC2 instance lifecycle events. Set the SQS queue as a target for the rule.
Answer: D (LEAVE A REPLY)
Amazon EC2 instances can send the state-change notification events to Amazon EventBridge. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-instance-state-changes.html Amazon EventBridge can send and receive events between event buses in AWS accounts. https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross-account.html

NEW QUESTION: 78
A developer supports an application that accesses data in an Amazon DynamoDB table. One of the item attributes is expirationDate in the timestamp format. The application uses this attribute to find items, archive them, and remove them from the table based on the timestamp value The application will be decommissioned soon, and the developer must find another way to implement this functionality. The developer needs a solution that will require the least amount of code to write.
Which solution will meet these requirements?
A. Enable TTL on the expirationDate attribute in the table. Create a DynamoDB stream. Create an AWS Lambda function to process the deleted items. Create a DynamoDB trigger for the Lambda function.
B. Create two AWS Lambda functions one to delete the items and one to process the items Create a DynamoDB stream Use the DeleteItem API operation to delete the items based on

the expirationDate attribute Use the GetRecords API operation to get the items from the DynamoDB stream and process them
C. Create two AWS Lambda functions, one to delete the items and one to process the items. Create an Amazon EventBndge scheduled rule to invoke the Lambda Functions Use the DeleteItem API operation to delete the items based on the expirationDate attribute. Use the GetRecords API operation to get the items from the DynamoDB table and process them.
D. Enable TTL on the expirationDate attribute in the table Specify an Amazon Simple Queue Service (Amazon SQS> dead-letter queue as the target to delete the items Create an AWS Lambda function to process the items
Answer: A (LEAVE A REPLY)
TTL for Automatic Deletion: DynamoDB's Time-to-Live effortlessly deletes expired items without manual intervention.
DynamoDB Stream: Captures changes to the table, including deletions of expired items, triggering downstream actions.
Lambda for Processing: A Lambda function connected to the stream provides custom logic for handling the deleted items.
Code Efficiency: This solution leverages native DynamoDB features and stream-based processing, minimizing the need for custom code.
Reference:
DynamoDB TTL Documentation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html DynamoDB Streams Documentation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html

NEW QUESTION: 79
A company is building a micro services app1 cation that consists of many AWS Lambda functions. The development team wants to use AWS Serverless Application Model (AWS SAM) templates to automatically test the Lambda functions. The development team plans to test a small percentage of traffic that is directed to new updates before the team commits to a full deployment of the application.
Which combination of steps will meet these requirements in the MOST operationally efficient way? (Select TWO.)
A. Use AWS SAM CLI commands in AWS CodeDeploy lo invoke the Lambda functions lo lest the deployment
B. Declare the EventInvokeConfig on the Lambda functions in the AWS SAM templates with OnSuccess and OnFailure configurations.
C. Enable gradual deployments through AWS SAM templates.
D. Set the deployment preference type to Canary10Percen130Minutes Use hooks to test the deployment.
E. Set the deployment preference type to Linear10PefcentEvery10Minutes Use hooks to test the deployment.

Answer: C,D (LEAVE A REPLY)

This solution will meet the requirements by using AWS Serverless Application Model (AWS SAM) templates and gradual deployments to automatically test the Lambda functions. AWS SAM templates are configuration files that define serverless applications and resources such as Lambda functions. Gradual deployments are a feature of AWS SAM that enable deploying new versions of Lambda functions incrementally, shifting traffic gradually, and performing validation tests during deployment. The developer can enable gradual deployments through AWS SAM templates by adding a DeploymentPreference property to each Lambda function resource in the template. The developer can set the deployment preference type to Canary10Percent30Minutes, which means that 10 percent of traffic will be shifted to the new version of the Lambda function for 30 minutes before shifting 100 percent of traffic. The developer can also use hooks to test the deployment, which are custom Lambda functions that run before or after traffic shifting and perform validation tests or rollback actions.

NEW QUESTION: 80

An application uses an Amazon EC2 Auto Scaling group. A developer notices that EC2 instances are taking a long time to become available during scale-out events. The UserData script is taking a long time to run.

The developer must implement a solution to decrease the time that elapses before an EC2 instance becomes available. The solution must make the most recent version of the application available at all times and must apply all available security updates. The solution also must minimize the number of images that are created. The images must be validated.

Which combination of steps should the developer take to meet these requirements? (Choose two.)

A. Use EC2 Image Builder to create an Amazon Machine Image (AMI). Install all the patches and agents that are needed to manage and run the application. Update the Auto Scaling group launch configuration to use the AMI.

B. Use EC2 Image Builder to create an Amazon Machine Image (AMI). Install the latest version of the application and all the patches and agents that are needed to manage and run the application. Update the Auto Scaling group launch configuration to use the AMI.

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

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

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

Answer: B,E (LEAVE A REPLY)

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

Set up AWS CodeDeploy to deploy the most recent version of the application at runtime. This will ensure that the application code is always up to date and does not depend on the AMI.

Remove any commands that perform operating system patching from the UserData script. This will reduce the time that the UserData script takes to run and speed up the instance launch process.
Reference:
[What Is AWS CloudFormation? - AWS CloudFormation]
[What Is AWS CodeDeploy? - AWS CodeDeploy]
[Running Commands on Your Linux Instance at Launch - Amazon Elastic Compute Cloud]

NEW QUESTION: 81
A company is building an application to accept data from customers. The data must be encrypted at rest and in transit.
The application uses an Amazon API Gateway API that resolves to AWS Lambda functions. The Lambda functions store the data in an Amazon Aurora MySQL DB cluster. The application worked properly during testing.
A developer configured an Amazon CloudFront distribution with field-level encryption that uses an AWS Key Management Service (AWS KMS) key. After the configuration of the distribution, the application behaved unexpectedly. All the data in the database changed from plaintext to ciphertext.
The developer must ensure that the data is not stored in the database as the ciphertext from the CloudFront field-level encryption.
Which solution will meet this requirement?
A. Change the CloudFront Viewer protocol policy from "HTTP and HTTPS" to "HTTPS only."
B. Add a Lambda function that uses the KMS key to decrypt the data fields before saving the data to the database.
C. Enable encryption on the DB cluster by using the same KMS key that is used in CloudFront.
D. Request and deploy a new SSL certificate to use with the CloudFront distribution.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 82
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 file. Create a new API Import the OpenAPI file Modify the new API to add request validation. Perform the tests Modify the existing API to add request validation. Deploy the existing API to production.
B. Modify the existing API to add request validation. Deploy the updated API to a new API Gateway stage Perform the tests Deploy the updated API to the API Gateway production stage.

C. Create a new API Add the necessary resources and methods including new request validation. Perform the tests Modify the existing API to add request validation. Deploy the existing API to production.
D. Clone the exiting API Modify the new API lo add request validation. Perform the tests Modify the existing API to add request validation Deploy the existing API to production.
Answer: (SHOW ANSWER)
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.

NEW QUESTION: 83
A developer is deploying a new application to Amazon Elastic Container Service (Amazon ECS). The developer needs to securely store and retrieve different types of variables. These variables include authentication information for a remote API, the URL for the API, and credentials. The authentication information and API URL must be available to all current and future deployed versions of the application across development, testing, and production environments.
How should the developer retrieve the variables with the FEWEST application changes?
A. Update the application to retrieve the variables from AWS Systems Manager Parameter Store. Use unique paths in Parameter Store for each variable in each environment. Store the credentials in AWS Secrets Manager in each environment.
B. Update the application to retrieve the variables from AWS Key Management Service (AWS KMS). Store the API URL and credentials as unique keys for each environment.
C. Update the application to retrieve the variables from an encrypted file that is stored with the application. Store the API URL and credentials in unique files for each environment.
D. Update the application to retrieve the variables from each of the deployed environments. Define the authentication information and API URL in the ECS task definition as unique names during the deployment process.
Answer: A (LEAVE A REPLY)
AWS Systems Manager Parameter Store is a service that provides secure, hierarchical storage for configuration data management and secrets management. The developer can update the application to retrieve the variables from Parameter Store by using the AWS SDK or the AWS CLI. The developer can use unique paths in Parameter Store for each variable in each environment, such as /dev/api-url, /test/api-url, and /prod/api-url. The developer can also store the credentials in AWS Secrets Manager, which is integrated with Parameter Store and provides additional features such as automatic rotation and encryption.
Reference:
[What Is AWS Systems Manager? - AWS Systems Manager]
[Parameter Store - AWS Systems Manager]

[What Is AWS Secrets Manager? - AWS Secrets Manager]

NEW QUESTION: 84
A company is creating a new application that gives users the ability to upload and share short video files. The average size of the video files is 10 MB. After a user uploads a file, a message needs to be placed into an Amazon Simple Queue Service (Amazon SQS) queue so the file can be processed. The files need to be accessible for processing within 5 minutes.
Which solution will meet these requirements MOST cost-effectively?
A. Write the files to Amazon S3 Glacier Deep Archive. Add the S3 location of the files to the SQS queue.
B. Write the files to Amazon S3 Standard. Add the S3 location of the files to the SQS queue.
C. Write the files to an Amazon Elastic Block Store (Amazon EBS) General Purpose SSD volume. Add the EBS location of the files to the SQS queue.
D. Write messages that contain the contents of the uploaded files to the SQS queue.
Answer: B (LEAVE A REPLY)
Comprehensive Detailed Explanation with all AWS Reference
Why Option B is Correct:
Amazon S3 Standard provides immediate access to files and is cost-effective for files that need to be accessed within 5 minutes.
By adding the S3 location to the SQS queue, you avoid transferring large files directly, which is both more efficient and scalable.
Why Other Options are Incorrect:
Option A: S3 Glacier Deep Archive is designed for archival storage with retrieval times ranging from minutes to hours, which does not meet the 5-minute requirement.
Option C: Amazon EBS is designed for block storage attached to EC2 instances, which adds unnecessary complexity and cost.
Option D: SQS is not designed to handle large file content directly and has message size limits (256 KB).
AWS Documentation Reference:
Amazon S3 Overview
Amazon SQS Best Practices

NEW QUESTION: 85
A developer is building a microservices-based application by using Python on AWS and several AWS services The developer must use AWS X-Ray The developer views the service map by using the console to view the service dependencies. During testing, the developer notices that some services are missing from the service map What can the developer do to ensure that all services appear in the X-Ray service map?
A. Modify the X-Ray Python agent configuration in each service to increase the sampling rate
B. Instrument the application by using the X-Ray SDK for Python. Install the X-Ray SDK for all the services that the application uses

C. Enable X-Ray data aggregation in Amazon CloudWatch Logs for all the services that the application uses
D. Increase the X-Ray service map timeout value in the X-Ray console
Answer: (SHOW ANSWER)
AWS X-Ray SDK: The primary way to enable X-Ray tracing within applications. The SDK sends data about requests and subsegments to the X-Ray daemon for service map generation.
Instrumenting All Services: To visualize a complete microservice architecture on the service map, each relevant service must include the X-Ray SDK.
Reference:
AWS X-Ray Documentation: https://docs.aws.amazon.com/xray/
X-Ray SDK for Python: https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-python.html

NEW QUESTION: 86
A company has an AWS Step Functions state machine named myStateMachine. The company configured a service role for Step Functions. The developer must ensure that only the myStateMachine state machine can assume the service role.
Which statement should the developer add to the trust policy to meet this requirement?
A. "Condition": { "ArnLike": { "aws:SourceArn":"urn:aws:states:ap-south-1:111111111111:stateMachine:myStateMachine" } }
B. "Condition": { "ArnLike": { "aws:SourceArn":"arn:aws:states:ap-south-1:*:stateMachine:myStateMachine" } }
C. "Condition": { "StringEquals": { "aws:SourceAccount": "111111111111" } }
D. "Condition": { "StringNotEquals": { "aws:SourceArn":"arn:aws:states:ap-south-1:111111111111:stateMachine:myStateMachine" } }
Answer: A (LEAVE A REPLY)
Comprehensive Detailed Explanation with all AWS Reference
Why Option A is Correct:
The ArnLike condition with the specific ARN for myStateMachine ensures that only this state machine can assume the role. The format urn:aws:states is correct for specifying Step Functions resources.
Why Other Options are Incorrect:
Option B: Wildcards (*) in the ARN allow more resources to assume the role, which violates the requirement.
Option C: This condition restricts the account but not the specific state machine.
Option D: A StringNotEquals condition is used to deny specific values, which does not ensure exclusivity for the desired state machine.
AWS Documentation Reference:
IAM Trust Policies for Step Functions

NEW QUESTION: 87
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. Configure a Lambda function destination with a failure condition Specify Lambda function as the destination type Specify the error-handling Lambda function's Amazon Resource Name (ARN) as the resource
B. Enable AWS X-Ray active tracing on the initial Lambda function. Configure X-Ray to capture stack traces of the failed invocations. Invoke the error-handling Lambda function by including the stack traces in the event object.
C. Configure a Lambda function trigger with a failure condition Specify Lambda function as the destination type Specify the error-handling Lambda function's Amazon Resource Name (ARN) as the resource
D. Create a status check alarm on the initial Lambda function. Configure the alarm to invoke the error-handling Lambda function when the alarm is initiated. Ensure that the alarm passes the stack trace in the event object.
Answer: A (LEAVE A REPLY)
Lambda Destinations on Failure: Allow routing asynchronous function invocations to specified resources (like another Lambda function) upon failure.
Error Handling: The error-handling Lambda receives details about the failure, enabling logging and custom actions.
Direct Integration: This solution leverages native Lambda functionality for a simpler implementation.

NEW QUESTION: 88
A company is providing read access to objects in an Amazon S3 bucket for different customers. The company uses 1AM permissions to restrict access to the S3 bucket. The customers can access only their own files.
Due to a regulation requirement, the company needs to enforce encryption in transit for interactions with Amazon S3.
Which solution will meet these requirements?
A. Add a bucket policy to the S3 bucket to deny S3 actions when the s3:x-amz-acl condition is equal to public-read.
B. Add an 1AM policy to the 1AM users that allows S3 actions when the s3:x-amz-acl condition is equal to bucket-owner-read.
C. Add an 1AM policy to the 1AM users to enforce the usage of the AWS SDK.
D. Add a bucket policy to the S3 bucket to deny S3 actions when the aws:SecureTransport condition is equal to false.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 89
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 system. Mount the EFS file system in Lambda. Store the result files and log file in the mount point. Append the log entries to the log file.
B. Create an Amazon Elastic Block Store (Amazon EBS) Multi-Attach enabled volume Attach the EBS volume to all Lambda functions. Update the Lambda function code to download the log file, append the log entries, and upload the modified log file to Amazon EBS
C. Create a reference to the /tmp local directory. Store the result files and log file by using the directory reference. Append the log entry to the log file.
D. 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 (LEAVE A REPLY)
Amazon EFS: A network file system (NFS) providing shared, scalable storage across multiple Lambda functions and other AWS resources.
Lambda Mounting: EFS file systems can be mounted within Lambda functions to access a shared storage space.
Log Appending: EFS supports appending data to existing files, making it ideal for the log file scenario.
Reference:
Amazon EFS Documentation: https://docs.aws.amazon.com/efs/
Using Amazon EFS with AWS Lambda:
https://docs.aws.amazon.com/lambda/latest/dg/services-efs.html

NEW QUESTION: 90
A developer previously deployed an AWS Lambda function as a .zip package. The developer needs to deploy the Lambda function as a container.
A. Create an Amazon ECR repository in the same AWS Region as the Lambda function. Package the Lambda function into a container image. Build the image and upload it to the Amazon ECR repository. Update the existing Lambda function configuration to specify the repository URI and container image tag.

B. Create an AWS SAM template that defines the Lambda function and its resources as code. Include a container image in the template, and store the container image in an Amazon S3 bucket. Deploy the AWS SAM template. Specify the S3 bucket URI.
C. Create an AWS CloudFormation template that defines the Lambda function and its resources as code. Include a container image in the template, and store the image in an Amazon S3 bucket. Deploy the CloudFormation template. Specify the S3 bucket URI.
D. Create an Amazon ECR repository in the same AWS Region as the Lambda function. Build the image and upload it to the Amazon ECR repository. Update the existing Lambda function to use the new image by specifying the repository URI.
Answer: (SHOW ANSWER)
Comprehensive Detailed Explanation with all AWS Reference
Why Option A is Correct:
Converting a Lambda function to use a container image involves packaging the function code into a container image, storing the image in Amazon Elastic Container Registry (ECR), and updating the function to use the ECR repository URI.
Why Other Options are Incorrect:
Option B: SAM templates support container-based Lambda deployment, but storing the image in S3 is not applicable.
Option C: CloudFormation does not natively support specifying Lambda container images in S3.
Option D: While partially correct, it omits the need to specify the image tag for the deployment.
AWS Documentation Reference:
Lambda Container Images

NEW QUESTION: 91
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 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.
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.

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.
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.
Answer: (SHOW ANSWER)
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.
Reference:
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

Valid DVA-C02 Dumps shared by TrainingQuiz.com for Helping Passing DVA-C02 Exam! TrainingQuiz.com now offer the **newest DVA-C02 exam dumps**, the TrainingQuiz.com DVA-C02 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C02 dumps with Test Engine here: https://www.trainingquiz.com/DVA-C02-practice-quiz.html (**649** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

NEW QUESTION: 92
A developer is receiving HTTP 400: ThrottlingException errors intermittently when calling the Amazon CloudWatch API. When a call fails, no data is retrieved.
What best practice should first be applied to address this issue?
A. Retry the call with exponential backoff.
B. Use the AWS CLI to get the metrics.
C. Contact AWS Support for a limit increase.
D. Analyze the applications and remove the API call.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 93
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 (LEAVE A REPLY)
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: 94
A team is developing an application that is deployed on Amazon EC2 instances. During testing, the team receives an error. The EC2 instances are unable to access an Amazon S3 bucket.
Which steps should the team take to troubleshoot this issue? (Select TWO.)
A. Check the S3 bucket policy to validate the access permissions for the S3 bucket.
B. Check the security groups that are assigned to the EC2 instances. Make sure that a rule is not blocking the access to Amazon S3.
C. Check whether the policy that is assigned to the 1AM user that is attached to the EC2 instances grants access to Amazon S3.
D. Check the S3 Lifecycle policy to validate the permissions that are assigned to the S3 bucket.
E. Check whether the policy that is assigned to the JAM role that is attached to the EC2 instances grants access to Amazon S3.
Answer: A,E (LEAVE A REPLY)

NEW QUESTION: 95
A developer is creating an AWS Serverless Application Model (AWS SAM) template. The AWS SAM template contains the definition of multiple AWS Lambda functions, an Amazon S3 bucket, and an Amazon CtoudFront distribution. One of the Lambda functions runs on Lambda@Edge in the CloudFront distribution. The S3 bucket is configured as an origin for the CloudFront distribution.
When the developer deploys the AWS SAM template in the eu-west-1 Region, the creation of the stack fails.
Which of the following could be the reason for this issue?
A. The CloudFront distribution and the S3 bucket cannot be created in the same Region.
B. Lambda@Edge functions can be created only in the us-east-1 Region.
C. CloudFront distributions can be created only in the us-east-1 Region.
D. A single AWS SAM template cannot contain multiple Lambda functions.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 96
A developer is incorporating AWS X-Ray into an application that handles personal identifiable information (PII). The application is hosted on Amazon EC2 instances. The application trace messages include encrypted PII and go to Amazon CloudWatch. The developer needs to ensure that no PII goes outside of the EC2 instances.
Which solution will meet these requirements?
A. Manually instrument the X-Ray SDK in the application code.

B. Use the X-Ray auto-instrumentation agent.
C. Use Amazon Macie to detect and hide PII. Call the X-Ray API from AWS Lambda.
D. Use AWS Distro for Open Telemetry.
Answer: A (LEAVE A REPLY)
This solution will meet the requirements by allowing the developer to control what data is sent to X-Ray and CloudWatch from the application code. The developer can filter out any PII from the trace messages before sending them to X-Ray and CloudWatch, ensuring that no PII goes outside of the EC2 instances. Option B is not optimal because it will automatically instrument all incoming and outgoing requests from the application, which may include PII in the trace messages. Option C is not optimal because it will require additional services and costs to use Amazon Macie and AWS Lambda, which may not be able to detect and hide all PII from the trace messages. Option D is not optimal because it will use Open Telemetry instead of X-Ray, which may not be compatible with CloudWatch and other AWS services.

NEW QUESTION: 97
A company is developing a serverless application that requires storage of sensitive API keys as environment variables for various services. The application requires the automatic rotation of the encryption keys every year.
Which solution will meet these requirements with no development effort?
A. Encrypt the environment variables by using AWS Secrets Manager. Set up automatic rotation in Secrets Manager.
B. Encrypt the environment variables by using AWS Systems Manager Parameter Store. Set up automatic rotation in Parameter Store.
C. Encrypt the environment variables by using AWS Key Management Service (AWS KMS) AWS managed keys. Configure a custom AWS Lambda function to automate key rotation.
D. Encrypt the environment variables by using AWS Key Management Service (AWS KMS) customer managed keys. Enable automatic key rotation.
Answer: (SHOW ANSWER)

NEW QUESTION: 98
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 instance. Store the unique identifier for each request in a database table. Modify the Lambda function to check the table for the identifier before processing the request.

B. Create an Amazon DynamoDB table. Store the unique identifier for each request in the table. Modify the Lambda function to check the table for the identifier before processing the request.
C. Create an Amazon DynamoDB table. Store the unique identifier for each request in the table. Modify the Lambda function to return a client error response when the function receives a duplicate request.
D. Create an Amazon ElastiCache for Memcached instance. Store the unique identifier for each request in the cache. Modify the Lambda function to check the cache for the identifier before processing the request.
Answer: (SHOW ANSWER)
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: 99
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 function. Configure VPC1 access for the function. Attach a security group named SG1 to both the Lambda function and the database. Configure the security group inbound and outbound rules to allow TCP traffic on Port 3306.
B. 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.
C. Create the Lambda function. Configure VPC1 access for the function. Assign a security group named SG1 to the Lambda function. Assign a second security group named SG2 to the database. Add an inbound rule to SG1 to allow TCP traffic from Port 3306.
D. 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 (LEAVE A REPLY)
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: 100
A developer is building a highly secure healthcare application using serverless components. This application requires writing temporary data to /Imp storage on an AWS Lambda function.
How should the developer encrypt this data?
A. Use OpenSSL to generate a symmetric encryption key on Lambda startup. Use this key to encrypt the data prior to writing to /tmp.
B. Enable Amazon EBS volume encryption with an AWS KMS key in the Lambda function configuration so that all storage attached to the Lambda function is encrypted.
C. Set up the Lambda function with a role and key policy to access an AWS KMS key. Use the key to generate a data key used to encrypt all data prior to writing to Amp storage.
D. Use an on-premises hardware security module (HSM) to generate keys, where the Lambda function requests a data key from the HSM and uses that to encrypt data on all requests to the function.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 101
A company runs a web application on Amazon EC2 instances behind an Application Load Balancer. The application uses Amazon DynamoDB as its database. The company wants to ensure high performance for reads and writes.
Which solution will meet these requirements MOST cost-effectively?
A. Configure auto-scaling for the DynamoDB table with a target utilization of 70%. Set the minimum and maximum capacity units based on the expected workload.
B. Use DynamoDB on-demand capacity mode for the table. Specify a maximum throughput higher than the expected peak read and write capacity units.
C. Use DynamoDB provisioned throughput mode for the table. Create an Amazon CloudWatch alarm on the ThrottledRequests metric. Invoke an AWS Lambda function to increase provisioned capacity.
D. Create an Amazon DynamoDB Accelerator (DAX) cluster. Configure the application to use the DAX endpoint.
Answer: (SHOW ANSWER)
Comprehensive Detailed Explanation with all AWS Reference
Why Option A is Correct:
Auto-scaling with a target utilization ensures the DynamoDB table dynamically adjusts capacity based on workload, maintaining high performance while optimizing cost. Setting a reasonable target utilization minimizes overprovisioning and throttling risks.
Why Other Options are Incorrect:

Option B: On-demand capacity is costlier than provisioned throughput for predictable workloads.
Option C: Using manual CloudWatch alarms and Lambda for scaling is less efficient and adds operational overhead.
Option D: DAX accelerates read performance but does not improve write performance.
AWS Documentation Reference:
DynamoDB Auto Scaling

NEW QUESTION: 102
A developer is creating an AWS Lambda function that needs network access to private resources in a VPC.
A. Attach the Lambda function to the VPC through private subnets. Create a security group that allows network access to the private resources. Associate the security group with the Lambda function.
B. Configure the Lambda function to route traffic through a VPN connection. Create a security group that allows network access to the private resources. Associate the security group with the Lambda function.
C. Configure a VPC endpoint connection for the Lambda function. Set up the VPC endpoint to route traffic through a NAT gateway.
D. Configure an AWS PrivateLink endpoint for the private resources. Configure the Lambda function to reference the PrivateLink endpoint.
Answer: A (LEAVE A REPLY)
Comprehensive Detailed Step by Step Explanation with All AWS Developer Reference:
When you need to provide an AWS Lambda function access to private resources in a VPC, the most common and straightforward approach is to attach the Lambda function to a VPC via private subnets. Once the Lambda function is associated with the VPC, you need to configure appropriate security groups to control the access to the private resources.
Lambda with VPC Access: Lambda functions can be attached to private subnets in a VPC, allowing them to access resources like RDS, EC2, or internal services within that VPC.
Security Groups: A security group acts as a virtual firewall for the Lambda function, ensuring that it can access only the necessary resources and ports in the VPC.
Alternatives:
Option B involves routing traffic through a VPN, which adds unnecessary complexity and operational overhead compared to simply attaching the Lambda to the VPC.
Option C requires configuring a VPC endpoint and a NAT gateway, which can be complex and costly.
Option D refers to AWS PrivateLink, which is used to access services over private connections, but it's unnecessary in this scenario unless you need a cross-VPC connection.
Reference:
Lambda functions in a VPC

NEW QUESTION: 103
A company has a multi-node Windows legacy application that runs on premises. The application uses a network shared folder as a centralized configuration repository to store configuration files in .xml format. The company is migrating the application to Amazon EC2 instances. As part of the migration to AWS, a developer must identify a solution that provides high availability for the repository.
Which solution will meet this requirement MOST cost-effectively?
A. Mount an Amazon Elastic Block Store (Amazon EBS) volume onto one of the EC2 instances. Deploy a file system on the EBS volume. Use the host operating system to share a folder. Update the application code to read and write configuration files from the shared folder.
B. Deploy a micro EC2 instance with an instance store volume. Use the host operating system to share a folder. Update the application code to read and write configuration files from the shared folder.
C. Create an Amazon S3 bucket to host the repository. Migrate the existing .xml files to the S3 bucket. Update the application code to use the AWS SDK to read and write configuration files from Amazon S3.
D. Create an Amazon S3 bucket to host the repository. Migrate the existing .xml files to the S3 bucket. Mount the S3 bucket to the EC2 instances as a local volume. Update the application code to read and write configuration files from the disk.
Answer: C (LEAVE A REPLY)
Amazon S3 is a service that provides highly scalable, durable, and secure object storage. The developer can create an S3 bucket to host the repository and migrate the existing .xml files to the S3 bucket. The developer can update the application code to use the AWS SDK to read and write configuration files from S3. This solution will meet the requirement of high availability for the repository in a cost-effective way.
Reference:
[Amazon Simple Storage Service (S3)]
[Using AWS SDKs with Amazon S3]

NEW QUESTION: 104
A company has an Amazon S3 bucket that contains sensitive data. The data must be encrypted in transit and at rest. The company encrypts the data in the S3 bucket by using an AWS Key Management Service (AWS KMS) key. A developer needs to grant several other AWS accounts the permission to use the S3 GetObject operation to retrieve the data from the S3 bucket.
How can the developer enforce that all requests to retrieve the data provide encryption in transit?
A. Define a resource-based policy on the S3 bucket to deny access when a request meets the condition "aws:SecureTransport": "false".
B. Define a resource-based policy on the S3 bucket to allow access when a request meets the condition "aws:SecureTransport": "false".

C. Define a role-based policy on the other accounts' roles to deny access when a request meets the condition of "aws:SecureTransport": "false".
D. Define a resource-based policy on the KMS key to deny access when a request meets the condition of "aws:SecureTransport": "false".
Answer: (SHOW ANSWER)
Amazon S3 supports resource-based policies, which are JSON documents that specify the permissions for accessing S3 resources. A resource-based policy can be used to enforce encryption in transit by denying access to requests that do not use HTTPS. The condition key aws:SecureTransport can be used to check if the request was sent using SSL. If the value of this key is false, the request is denied; otherwise, the request is allowed. Reference: How do I use an S3 bucket policy to require requests to use Secure Socket Layer (SSL)?

NEW QUESTION: 105
A company has an application that is deployed on AWS Elastic Beanstalk. The application generates user-specific PDFs and stores the PDFs in an Amazon S3 bucket. The application then uses Amazon Simple Email Service (Amazon SES) to send the PDFs by email to subscribers.
Users no longer access the PDFs 90 days after the PDFs are generated. The S3 bucket is not versioned and contains many obsolete PDFs.
A developer must reduce the number of files in the S3 bucket by removing PDFs that are older than 90 days.
Which solution will meet this requirement with the LEAST development effort?
A. Partition the S3 objects with a <year>/<month>/<day> key prefix. Create an AWS Lambda function to remove objects that have prefixes that have reached the expiration date.
B. Create an AWS Lambda function. Program the Lambda function to scan all the objects in the S3 bucket every day and to delete objects after 90 days.
C. Create an S3 Lifecycle rule for the S3 bucket to expire objects after 90 days.
D. Update the application code. In the code, add a rule to scan all the objects in the S3 bucket every day and to delete objects after 90 days.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 106
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 notification. Add an Amazon API Gateway API to invoke the function. Call the API from the client side when login confirmation is received.

B. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification. Add an Amazon Cognito post authentication Lambda trigger for the function.
C. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification. Create an Amazon CloudWatch Logs log subscription filter to invoke the function based on the login status.
D. Configure Amazon Cognito to stream all logs to Amazon Kinesis Data Firehose. 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 (LEAVE A REPLY)
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

Valid DVA-C02 Dumps shared by TrainingQuiz.com for Helping Passing DVA-C02 Exam! TrainingQuiz.com now offer the **newest DVA-C02 exam dumps**, the TrainingQuiz.com DVA-C02 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C02 dumps with Test Engine here: https://www.trainingquiz.com/DVA-C02-practice-quiz.html (**649** Q&As Dumps, **40%OFF**
Special Discount: **Exam-Tests**)

NEW QUESTION: 107
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 (LEAVE A REPLY)
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.

NEW QUESTION: 108
A developer is updating an Amazon API Gateway REST API to have a mock endpoint. The developer wants to update the integration request mapping template so the endpoint will respond to mock integration requests with specific HTTP status codes based on various conditions.
A. { if($input.params('integration') == "mock") "statusCode": 404 else "statusCode": 500 end }
B. { if($input.params('scope') == "internal") "statusCode": 200 else "statusCode": 500 end }
C. { if($input.path("integration")) "statusCode": 200 else "statusCode":404 end }
D. { if($context.integration.status) "statusCode": 200 else "statusCode": 500 end }
Answer: (SHOW ANSWER)
Comprehensive Detailed Step by Step Explanation with All AWS Developer Reference:
In this scenario, the developer is configuring a mock integration in API Gateway. The integration request mapping template allows you to map the incoming request data to a format that the API expects. For mock integration, it's common to return specific HTTP status codes based on the conditions.
Using $context.integration.status: The $context.integration.status variable refers to the status of the API Gateway integration, which is useful for generating responses based on the condition. Option D correctly uses this variable to determine the HTTP status code, returning 200 for a successful mock request or 500 for a failure.
Alternatives:
Options A, B, and C do not use the correct context variables for handling mock integrations. These options would not return the correct status codes based on the actual integration status.
Reference:
API Gateway Mapping Templates and Accessing Context Variables

NEW QUESTION: 109
A company needs to package and deploy an application that uses AWS Lambda to compress and decompress video clips. The application uses a video codec library that is larger than 250 MB. The application uses the library to compress the videos before storage and to decompress the videos upon retrieval.
A. Create one Lambda function. Upload one zip file that contains code to handle video compression and decompression to the function. Include the codec library in the zip file.
B. Create two Lambda functions. Upload one zip file that contains code to handle video compression to one function. Upload a second zip file that contains code for video decompression to the second function. Include the codec library in both zip files.
C. Create two Lambda functions. Upload one zip file that contains code to handle video compression to one function. Upload a second zip file that contains code for video

decompression to the second function. Create one Lambda layer for the codec library. Add the layer to both functions.

D. Create two Lambda functions. Build one container image that contains code to handle video compression and a second image that contains video decompression code. Add the codec library to both images. Upload the images to Amazon ECR. Use the containers to create the Lambda functions.

Answer: (SHOW ANSWER)

Comprehensive and Detailed Step-by-Step

Option D: Use Lambda with Container Images

AWS Lambda supports container images up to 10 GB in size, making it suitable for applications with large dependencies, such as a video codec library larger than 250 MB.

By creating separate container images for video compression and decompression, the application can efficiently isolate functionality while ensuring that each function includes the required dependencies.

The container images are stored in Amazon ECR and used to create the Lambda functions.

Why Other Options Are Incorrect:

Option A: A single Lambda function with all functionalities and dependencies in one zip file is not feasible due to the 250 MB deployment package size limit for zip files.

Option B: Including the library in two separate zip files still exceeds the size limit for Lambda zip deployment packages.

Option C: While using a Lambda layer can reduce redundancy, the combined size of the layer and the zip files would exceed the limit of 250 MB.

Reference:

Using Container Images with AWS Lambda

NEW QUESTION: 110

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

Which solution will meet these requirements?

A. Create API Gateway resources and set the integration type value to MOCK Configure the method integration request and integration response to associate a response with an HTTP status code Create an API Gateway stage and deploy the API.

B. Create an AWS Lambda function that returns mocked responses and various HTTP status codes. Create API Gateway resources and set the integration type value to AWS_PROXY Deploy the API.

C. Create an EC2 application that returns mocked HTTP responses Create API Gateway resources and set the integration type value to AWS Create an API Gateway stage and deploy the API.

D. Create API Gateway resources and set the integration type value set to HTTP_PROXY. Add mapping templates and deploy the API. Create an AWS Lambda layer that returns various HTTP status codes Associate the Lambda layer with the API deployment

Answer: A (LEAVE A REPLY)

API Gateway Mocking: This feature is built for decoupling development dependencies. Here's the process:

Create resources and methods in your API Gateway.

Set the integration type to 'MOCK'.

Define Integration Responses, mapping HTTP status codes to desired mocked responses (JSON, etc.).

Deployment and Use:

Create a deployment stage for the API.

Frontend teams can call this API and get the mocked responses without a real backend.

Reference:

Mocking API Gateway APIs:

https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-mock-integration.html

NEW QUESTION: 111

A developer is troubleshooting an application in an integration environment. In the application, an Amazon Simple Queue Service (Amazon SQS) queue consumes messages and then an AWS Lambda function processes the messages. The Lambda function transforms the messages and makes an API call to a third-party service.

There has been an increase in application usage. The third-party API frequently returns an HTTP 429 Too Many Requests error message. The error message prevents a significant number of messages from being processed successfully.

How can the developer resolve this issue?

A. Increase the SQS event source's batch size setting.

B. Configure provisioned concurrency for the Lambda function based on the third-party API's documented rate limits.

C. Increase the retry attempts and maximum event age in the Lambda function's asynchronous configuration.

D. Configure maximum concurrency on the SQS event source based on the third-party service's documented rate limits.

Answer: D (LEAVE A REPLY)

Maximum concurrency for SQS as an event source allows customers to control the maximum concurrent invokes by the SQS event source1. When multiple SQS event sources are configured to a function, customers can control the maximum concurrent invokes of individual SQS event source1.

In this scenario, the developer needs to resolve the issue of the third-party API frequently returning an HTTP 429 Too Many Requests error message, which prevents a significant

number of messages from being processed successfully. To achieve this, the developer can follow these steps:
Find out the documented rate limits of the third-party API, which specify how many requests can be made in a given time period.
Configure maximum concurrency on the SQS event source based on the rate limits of the third-party API. This will limit the number of concurrent invokes by the SQS event source and prevent exceeding the rate limits of the third-party API.
Test and monitor the application performance and adjust the maximum concurrency value as needed.
By using this solution, the developer can reduce the frequency of HTTP 429 errors and improve the message processing success rate. The developer can also avoid throttling or blocking by the third-party API.

NEW QUESTION: 112
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: (SHOW ANSWER)
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.

NEW QUESTION: 113
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 Gateway. Configure Amazon CloudWatch Logs as the delivery stream's destination.
B. Turn on AWS CloudTrail Insights and create a trail Specify the Amazon Resource Name (ARN) of the trail for the stage of the API.
C. 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.

D. Turn on execution logging and access logging in Amazon CloudWatch Logs for the API stage. Create a CloudWatch Logs log group. Specify the Amazon Resource Name (ARN) of the log group for the API stage.

Answer: D (LEAVE A REPLY)

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.

NEW QUESTION: 114

A developer has written a distributed application that uses micro services. The microservices are running on Amazon EC2 instances. Because of message volume, the developer is unable to match log output from each microservice to a specific transaction. The developer needs to analyze the message flow to debug the application.

Which combination of steps should the developer take to meet this requirement? (Select TWO.)

A. Enable AWS X-Ray. Configure Amazon CloudWatch to push logs to X-Ray.

B. Configure an interface VPC endpoint to allow traffic to reach the global AWS X-Ray daemon on TCP port 2000.

C. Set up Amazon CloudWatch metric streams to collect streaming data from the microservices.

D. Add the AWS X-Ray software development kit (SDK) to the microservices. Use X-Ray to trace requests that each microservice makes.

E. Download the AWS X-Ray daemon. Install the daemon on an EC2 instance. Ensure that the EC2 instance allows UDP traffic on port 2000.

Answer: D,E (LEAVE A REPLY)

NEW QUESTION: 115

A developer is planning to migrate on-premises company data to Amazon S3. The data must be encrypted, and the encryption Keys must support automate annual rotation. The company must use AWS Key Management Service (AWS KMS) to encrypt the data.
When type of keys should the developer use to meet these requirements?
A. Amazon S3 managed keys
B. Symmetric customer managed keys with key material that is generated by AWS
C. Asymmetric customer managed keys with key material that generated by AWS
D. Symmetric customer managed keys with imported key material
Answer: B (LEAVE A REPLY)
The type of keys that the developer should use to meet the requirements is symmetric customer managed keys with key material that is generated by AWS. This way, the developer can use AWS Key Management Service (AWS KMS) to encrypt the data with a symmetric key that is managed by the developer. The developer can also enable automatic annual rotation for the key, which creates new key material for the key every year. The other options either involve using Amazon S3 managed keys, which do not support automatic annual rotation, or using asymmetric keys or imported key material, which are not supported by S3 encryption.

NEW QUESTION: 116
A company uses AWS X-Ray to monitor a serverless application. The components of the application have different request rates. The user interactions and transactions are important to trace, but they are low in volume. The background processes such as application health checks, polling, and connection maintenance generate high volumes of read-only requests. Currently, the default X-Ray sampling rules are universal for all requests. Only the first request per second and some additional requests are recorded. This setup is not helping the company review the requests based on service or request type.
A developer must configure rules to trace requests based on service or request properties. The developer must trace the user interactions and transactions without wasting effort recording minor background tasks.
Which solution will meet these requirements?
A. Disable sampling and trace all requests for requests that handle user interactions or transactions. Sample high-volume read-only requests at a lower rate.
B. Disable sampling for high-volume read-only requests. Sample at a higher rate for all requests that handle user interactions or transactions.
C. Disable sampling and trace all requests for requests that handle user interactions or transactions. Sample high-volume read-only requests at a higher rate.
D. Disable sampling for high-volume read-only requests. Sample at a lower rate for all requests that handle user interactions or transactions.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 117

A company runs a new application on AWS Elastic Beanstalk. The company needs to deploy updates to the application. The updates must not cause any downtime for application users. The deployment must forward a specified percentage of incoming client traffic to a new application version during an evaluation period.
Which deployment type will meet these requirements?
A. Rolling
B. Traffic-splitting
C. In-place
D. Immutable
Answer: B (LEAVE A REPLY)
AWS Elastic Beanstalk supports several deployment policies, and in this case, the requirement is to forward a specific percentage of traffic to the new version without causing downtime. The Traffic-splitting deployment policy is the most appropriate choice.
Traffic-splitting Deployment: This deployment method allows you to gradually shift a specified percentage of incoming traffic from the old environment version to the new one. During the evaluation period, if any issues are detected, the traffic can be redirected back to the old version.
No Downtime: This method ensures no downtime since both versions of the application run concurrently, and traffic is split between them.
Alternatives:
Rolling deployments (Option A): These gradually replace instances but may result in partial downtime if some instances fail during deployment.
In-place deployments (Option C): In-place deployments replace instances without creating new ones, which can lead to downtime.
Immutable deployments (Option D): While this ensures no downtime by creating entirely new instances, it doesn't provide traffic splitting during the evaluation phase.
Reference:
Elastic Beanstalk Deployment Policies

NEW QUESTION: 118
A company has an Amazon S3 bucket containing premier content that it intends to make available to only paid subscribers of its website. The S3 bucket currently has default permissions of all objects being private to prevent inadvertent exposure of the premier content to non-paying website visitors.
How can the company Limit the ability to download a premier content file in the S3 Bucket to paid subscribers only?
A. Apply a bucket policy that allows anonymous users to download the content from the S3 bucket.
B. Generate a pre-signed object URL for the premier content file when a pad subscriber requests a download.

C. Add a Docket policy that requires multi-factor authentication for request to access the S3 bucket objects.
D. Enable server-side encryption on the S3 bucket for data protection against the non-paying website visitors.
Answer: (SHOW ANSWER)
This solution will limit the ability to download a premier content file in the S3 bucket to paid subscribers only because it uses a pre-signed object URL that grants temporary access to an S3 object for a specified duration. The pre-signed object URL can be generated by the company's website when a paid subscriber requests a download, and can be verified by Amazon S3 using the signature in the URL. Option A is not optimal because it will allow anyone to download the content from the S3 bucket without verifying their subscription status. Option C is not optimal because it will require additional steps and costs to configure multi-factor authentication for accessing the S3 bucket objects, which may not be feasible or user-friendly for paid subscribers. Option D is not optimal because it will not prevent non-paying website visitors from accessing the S3 bucket objects, but only encrypt them at rest.

NEW QUESTION: 119
A developer creates an AWS Lambda function that is written in Java. During testing, the Lambda function does not work how the developer expected. The developer wants to use tracing capabilities to troubleshoot the problem.
Which AWS service should the developer use to accomplish this goal?
A. AWS Trusted Advisor
B. AWS X-Ray
C. AWS CloudTrail
D. Amazon CloudWatch
Answer: B (LEAVE A REPLY)

NEW QUESTION: 120
A team of developed is using an AWS CodePipeline pipeline as a continuous integration and continuous delivery (CI/CD) mechanism for a web application. A developer has written unit tests to programmatically test the functionality of the application code. The unit tests produce a test report that shows the results of each individual check. The developer now wants to run these tests automatically during the CI/CD process.
A. Write a Git pre-commit hook that runs the test before every commit. Ensure that each developer who is working on the project has the pre-commit hook instated locally. Review the test report and resolve any issues before pushing changes to AWS CodeCommit.
B. Add a new stage to the pipeline. Use AWS CodeBuild as the provider. Add the new stage after the stage that deploys code revisions to the test environment. Write a buildspec that fails the CodeBuild stage if any test does not pass. Use the test reports feature of Codebuild to integrate the report with the CodoBuild console. View the test results in CodeBuild Resolve any issues.

C. Add a new stage to the pipeline. Use AWS CodeBuild at the provider. Add the new stage before the stage that deploys code revisions to the test environment. Write a buildspec that fails the CodeBuild stage it any test does not pass. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild console. View the test results in codeBuild Resolve any issues.
D. Add a new stage to the pipeline. Use Jenkins as the provider. Configure CodePipeline to use Jenkins to run the unit tests. Write a Jenkinsfile that fails the stage if any test does not pass. Use the test report plugin for Jenkins to integrate the repot with the Jenkins dashboard. View the test results in Jenkins. Resolve any issues.
Answer: (SHOW ANSWER)
The solution that will meet the requirements is to add a new stage to the pipeline. Use AWS CodeBuild as the provider. Add the new stage before the stage that deploys code revisions to the test environment. Write a buildspec that fails the CodeBuild stage if any test does not pass. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild console. View the test results in CodeBuild. Resolve any issues. This way, the developer can run the unit tests automatically during the CI/CD process and catch any bugs before deploying to the test environment. The developer can also use the test reports feature of CodeBuild to view and analyze the test results in a graphical interface. The other options either involve running the tests manually, running them after deployment, or using a different provider that requires additional configuration and integration.

NEW QUESTION: 121
A developer needs to use a code template to create an automated deployment of an application onto Amazon EC2 instances. The template must be configured to repeat deployment, installation, and updates of resources for the application. The template must be able to create identical environments and roll back to previous versions.
Which solution will meet these requirements?
A. Use AWS Amplify for automatic deployment templates. Use a traffic-splitting deployment to copy any deployments. Modify any resources created by Amplify, if necessary.
B. Use AWS CodeBuild for automatic deployment. Upload the required AppSpec file template. Save the appspec.yml file in the root directory folder of the revision. Specify the deployment group that includes the EC2 instances for the deployment.
C. Use AWS CloudFormation to create an infrastructure template in JSON format to deploy the EC2 instances. Use Cloud Formation helper scripts to install the necessary software and to start the application. Call the scripts directly from the template.
D. Use AWS AppSync to deploy the application. Upload the template as a GraphQL schema. Specify the EC2 instances for deployment of the application. Use resolvers as a version control mechanism and to make any updates to the deployments.
Answer: (SHOW ANSWER)

Valid DVA-C02 Dumps shared by TrainingQuiz.com for Helping Passing DVA-C02 Exam! TrainingQuiz.com now offer the **newest DVA-C02 exam dumps**, the TrainingQuiz.com DVA-C02 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C02 dumps with Test Engine here: https://www.trainingquiz.com/DVA-C02-practice-quiz.html (**649** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

NEW QUESTION: 122
A company wants to automate part of its deployment process. A developer needs to automate the process of checking for and deleting unused resources that supported previously deployed stacks but that are no longer used.
The company has a central application that uses the AWS Cloud Development Kit (AWS CDK) to manage all deployment stacks. The stacks are spread out across multiple accounts. The developer's solution must integrate as seamlessly as possible within the current deployment process.
Which solution will meet these requirements with the LEAST amount of configuration?
A. In the central AWS CDK application, write a handler function in the code that uses AWS SDK calls to check for and delete unused resources. Create an AWS CloudPormation template from a JSON file. Use the template to attach the function code to an AWS Lambda function and lo invoke the Lambda function when the deployment slack runs.
B. In the central AWS CDK application. write a handler function in the code that uses AWS SDK calls to check for and delete unused resources. Create an AWS CDK custom resource Use the custom resource to attach the function code to an AWS Lambda function and to invoke the Lambda function when the deployment stack runs.
C. In the central AWS CDK, write a handler function m the code that uses AWS SDK calls to check for and delete unused resources. Create an API in AWS Amplify Use the API to attach the function code to an AWS Lambda function and to invoke the Lambda function when the deployment stack runs.
D. In the AWS Lambda console write a handler function in the code that uses AWS SDK calls to check for and delete unused resources. Create an AWS CDK custom resource. Use the custom resource to import the Lambda function into the stack and to Invoke the Lambda function when the deployment stack runs.
Answer: B (LEAVE A REPLY)
This solution meets the requirements with the least amount of configuration because it uses a feature of AWS CDK that allows custom logic to be executed during stack deployment or deletion. The AWS Cloud Development Kit (AWS CDK) is a software development framework that allows you to define cloud infrastructure as code and provision it through CloudFormation. An AWS CDK custom resource is a construct that enables you to create resources that are not natively supported by CloudFormation or perform tasks that are not supported by CloudFormation during stack deployment or deletion. The developer can write a handler

function in the code that uses AWS SDK calls to check for and delete unused resources, and create an AWS CDK custom resource that attaches the function code to a Lambda function and invokes it when the deployment stack runs. This way, the developer can automate the cleanup process without requiring additional configuration or integration. Creating a CloudFormation template from a JSON file will require additional configuration and integration with the central AWS CDK application. Creating an API in AWS Amplify will require additional configuration and integration with the central AWS CDK application and may not provide optimal performance or availability. Writing a handler function in the AWS Lambda console will require additional configuration and integration with the central AWS CDK application.

NEW QUESTION: 123
A company has implemented a pipeline in AWS CodePipeline. The company Is using a single AWS account and does not use AWS Organizations. The company needs to test its AWS CloudFormation templates in its primary AWS Region and a disaster recovery Region.
Which solution will meet these requirements with the MOST operational efficiency?
A. Use the Snyk action in CodePipeline to deploy and test the CloudFormation templates in each Region.
B. Configure CodePipeline to invoke AWS CodeBuild to deploy and test the CloudFormation templates in each Region. Update CodeBuild and CloudFormation with appropriate permissions.
C. In the CodePipeline pipeline, implement an AWS CodeDeploy action for each Region to deploy and test the Cloud Formation templates. Update CodePipeline and AWS CodeBuild with appropriate permissions.
D. Configure CodePipeline to deploy and test the Cloud Formation templates. Use CloudFormation StackSets to start deployment across both Regions.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 124
A developer is investigating an issue in part of a company's application. In the application messages are sent to an Amazon Simple Queue Service (Amazon SQS) queue The AWS Lambda function polls messages from the SQS queue and sends email messages by using Amazon Simple Email Service (Amazon SES) Users have been receiving duplicate email messages during periods of high traffic.
Which reasons could explain the duplicate email messages? (Select TWO.)
A. Standard SQS queues support at-least-once message delivery
B. Standard SQS queues support exactly-once processing, so the duplicate email messages are because of user error.
C. Amazon SES has the DomainKeys Identified Mail (DKIM) authentication incorrectly configured
D. The SQS queue's visibility timeout is lower than or the same as the Lambda function's timeout.

E. The Amazon SES bounce rate metric is too high.

Answer: (SHOW ANSWER)

SQS Delivery Behavior: Standard SQS queues guarantee at-least-once delivery, meaning messages may be processed more than once. This can lead to duplicate emails in this scenario.

Visibility Timeout: If the visibility timeout on the SQS queue is too short, a message might become visible for another consumer before the first Lambda function finishes processing it. This can also lead to duplicates.

Reference:

Amazon SQS Delivery Semantics: [invalid URL removed]

Amazon SQS Visibility Timeout:

https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html

NEW QUESTION: 125

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 (LEAVE A REPLY)

CodeDeploy Predefined Configurations: CodeDeploy offers built-in deployment configurations for common scenarios.

Canary Deployment: Canary deployments gradually shift traffic to a new version, ideal for controlled rollouts like this requirement.

CodeDeployDefault.ECSCanary10Percent15Minutes: This configuration matches the company's requirements, shifting 10% of traffic initially and then completing the rollout after 15 minutes.

Reference:

AWS CodeDeploy Deployment Configurations:

https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations-create.html

NEW QUESTION: 126

A company runs an application on AWS The application stores data in an Amazon DynamoDB table Some queries are taking a long time to run These slow queries involve an attribute that is

not the table's partition key or sort key The amount of data that the application stores in the DynamoDB table is expected to increase significantly. A developer must increase the performance of the queries.
Which solution will meet these requirements'?
A. Increase the page size for each request by setting the Limit parameter to be higher than the default value Configure the application to retry any request that exceeds the provisioned throughput.
B. Create a global secondary index (GSI). Set query attribute to be the partition key of the index
C. Perform a parallel scan operation by issuing individual scan requests in the parameters specify the segment for the scan requests and the total number of segments for the parallel scan.
D. Turn on read capacity auto scaling for the DynamoDB table. Increase the maximum read capacity units (RCUs).
Answer: (SHOW ANSWER)
Global Secondary Index (GSI): GSIs enable alternative query patterns on a DynamoDB table by using different partition and sort keys.
Addressing Query Bottleneck: By making the slow-query attribute the GSI's partition key, you optimize queries on that attribute.
Scalability: GSIs automatically scale to handle increasing data volumes.
Reference:
Amazon DynamoDB Global Secondary Indexes:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GSI.html

NEW QUESTION: 127
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 (LEAVE A REPLY)
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.

NEW QUESTION: 128
A web application is using Amazon Kinesis Data Streams for clickstream data that may not be consumed for up to 12 hours.
How can the developer implement encryption at rest for data within the Kinesis Data Streams?
A. Use Amazon Kinesis Consumer Library.
B. Enable SSL connections to Kinesis.
C. Enable server-side encryption in Kinesis Data Streams.
D. Encrypt the data once it is at rest with a Lambda function.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 129
A company developed an API application on AWS by using Amazon CloudFront. Amazon API Gateway, and AWS Lambda. The API has a minimum of four requests every second A developer notices that many API users run the same query by using the POST method. The developer wants to cache the POST request to optimize the API resources.
Which solution will meet these requirements'?
A. Configure the CloudFront cache Update the application to return cached content based upon the default request headers.
B. Override the cache method in me selected stage of API Gateway Select the POST method.
C. Save the latest request response in Lambda /tmp directory Update the Lambda function to check the /tmp directory
D. Save the latest request m AWS Systems Manager Parameter Store Modify the Lambda function to take the latest request response from Parameter Store
Answer: A (LEAVE A REPLY)
This solution will meet the requirements by using Amazon CloudFront, which is a content delivery network (CDN) service that speeds up the delivery of web content and APIs to end users. The developer can configure the CloudFront cache, which is a set of edge locations that store copies of popular or recently accessed content close to the viewers. The developer can also update the application to return cached content based upon the default request headers, which are a set of HTTP headers that CloudFront automatically forwards to the origin server and uses to determine whether an object in an edge location is still valid. By caching the POST requests, the developer can optimize the API resources and reduce the latency for repeated

queries. Option B is not optimal because it will override the cache method in the selected stage of API Gateway, which is not possible or effective as API Gateway does not support caching for POST methods by default. Option C is not optimal because it will save the latest request response in Lambda /tmp directory, which is a local storage space that is available for each Lambda function invocation, not a cache that can be shared across multiple invocations or requests. Option D is not optimal because it will save the latest request in AWS Systems Manager Parameter Store, which is a service that provides secure and scalable storage for configuration data and secrets, not a cache for API responses.

NEW QUESTION: 130
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 (LEAVE A REPLY)
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.

NEW QUESTION: 131
A developer has created an AWS Lambda function to provide notification through Amazon Simple Notification Service (Amazon SNS) whenever a file is uploaded to Amazon S3 that is larger than 50 MB. The developer has deployed and tested the Lambda function by using the CLI. However, when the event notification is added to the S3 bucket and a 3.000 MB file is uploaded, the Lambda function does not launch.

Which of the following Is a possible reason for the Lambda function's inability to launch?
A. Lambda functions cannot be invoked directly from an S3 event.
B. The resource-based policy for the Lambda function does not have the required permissions to be invoked by Amazon S3.
C. The S3 bucket needs to be made public.
D. The S3 event notification does not activate for files that are larger than 1.000 MB.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 132
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 Region. Enable secret replication to the secondary Region Update the application to use the Amazon Resource Name (ARN) based on the Region.
B. Store credentials in AWS Systems Manager Parameter Store in the primary Region. Enable parameter replication to the secondary Region. Update the application to use the Amazon Resource Name (ARN) based on the Region.
C. Store credentials in a config file. Upload the config file to an S3 bucket in me primary Region. Enable Cross-Region Replication (CRR) to an S3 bucket in the secondary region. Update the application to access the config file from the S3 bucket based on the Region.
D. Store credentials in a config file. Upload the config file to an Amazon Elastic File System (Amazon EFS) file system. 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 (LEAVE A REPLY)
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.
Reference:
AWS Secrets Manager
Replicating and sharing secrets
Using your own encryption keys

NEW QUESTION: 133
A developer has a legacy application that is hosted on-premises. Other applications hosted on AWS depend on the on-premises application for proper functioning. In case of any application errors, the developer wants to be able to use Amazon CloudWatch to monitor and troubleshoot all applications from one place.
How can the developer accomplish this?
A. Install an AWS SDK on the on-premises server to automatically send logs to CloudWatch.
B. Download the CloudWatch agent to the on-premises server. Configure the agent to use IAM user credentials with permissions for CloudWatch.
C. Upload log files from the on-premises server to Amazon S3 and have CloudWatch read the files.
D. Upload log files from the on-premises server to an Amazon EC2 instance and have the instance forward the logs to CloudWatch.
Answer: B (LEAVE A REPLY)
Amazon CloudWatch is a service that monitors AWS resources and applications. The developer can use CloudWatch to monitor and troubleshoot all applications from one place. To do so, the developer needs to download the CloudWatch agent to the on-premises server and configure the agent to use IAM user credentials with permissions for CloudWatch. The agent will collect logs and metrics from the on-premises server and send them to CloudWatch.
Reference:
[What Is Amazon CloudWatch? - Amazon CloudWatch]
[Installing and Configuring the CloudWatch Agent - Amazon CloudWatch]

NEW QUESTION: 134
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 (LEAVE A REPLY)

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.

NEW QUESTION: 135

A developer is creating an application that uses an Amazon DynamoDB table. The developer needs to develop code that reads all records that were added to the table during the previous day. creates HTML reports, and pushes the reports into third-party storage. The item size varies from 1 KB to 4 KB, and the index structure is defined with the date. The developer needs to minimize the read capacity that the application requires from the DynamoDB table.

Which DynamoDB API operation should the developer use in the code to meet these requirements?

A. BatchGetItem

B. GetItem

C. Query

D. Scan

Answer: C (LEAVE A REPLY)

NEW QUESTION: 136

A company receives food orders from multiple partners. The company has a microservices application that uses Amazon API Gateway APIs with AWS Lambda integration. Each partner sends orders by calling a customized API that is exposed through API Gateway. The API call invokes a shared Lambda function to process the orders.

Partners need to be notified after the Lambda function processes the orders. Each partner must receive updates for only the partner's own orders. The company wants to add new partners in the future with the fewest code changes possible.

Which solution will meet these requirements in the MOST scalable way?

A. Create a different Amazon Simple Notification Service (Amazon SNS) topic for each partner. Configure the Lambda function to publish messages for each partner to the partner's SNS topic.
B. Create a different Lambda function for each partner. Configure the Lambda function to notify each partner's service endpoint directly.
C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Configure the Lambda function to publish messages with specific attributes to the SNS topic. Subscribe each partner to the SNS topic. Apply the appropriate filter policy to the topic subscriptions.
D. Create one Amazon Simple Notification Service (Amazon SNS) topic. Subscribe all partners to the SNS topic.
Answer: C (LEAVE A REPLY)
Amazon Simple Notification Service (Amazon SNS) is a fully managed messaging service that enables pub/sub communication between distributed systems. The developer can create an SNS topic and configure the Lambda function to publish messages with specific attributes to the topic. The developer can subscribe each partner to the SNS topic and apply the appropriate filter policy to the topic subscriptions. This way, each partner will receive updates for only their own orders based on the message attributes. This solution will meet the requirements in the most scalable way and allow adding new partners in the future with minimal code changes.
Reference:
[Amazon Simple Notification Service (SNS)]
[Filtering Messages with Attributes - Amazon Simple Notification Service]

Valid DVA-C02 Dumps shared by TrainingQuiz.com for Helping Passing DVA-C02 Exam! TrainingQuiz.com now offer the **newest DVA-C02 exam dumps**, the TrainingQuiz.com DVA-C02 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C02 dumps with Test Engine here: https://www.trainingquiz.com/DVA-C02-practice-quiz.html (**649** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

NEW QUESTION: 137
A company's application has an AWS Lambda function that processes messages from IoT devices. The company wants to monitor the Lambda function to ensure that the Lambda function is meeting its required service level agreement (SLA).
A developer must implement a solution to determine the application's throughput in near real time. The throughput must be based on the number of messages that the Lambda function receives and processes in a given time period. The Lambda function performs initialization and post-processing steps that must not factor into the throughput measurement.
What should the developer do to meet these requirements?

A. Modify the application to log the calculated throughput to Amazon CloudWatch Logs. Use Amazon EventBridge to invoke a separate Lambda function to process the logs on a schedule.
B. Use the Lambda function's Invocations metric and Duration metric to calculate the throughput in Amazon CloudWatch.
C. Modify the application to publish custom Amazon CloudWatch metrics when the Lambda function receives and processes each message. Use the metrics to calculate the throughput.
D. Use the Lambda function's ConcurrentExecutions metric in Amazon CloudWatch to measure the throughput.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 138
A developer is writing an AWS Lambda function. The developer wants to log key events that occur while the Lambda function runs. The developer wants to include a unique identifier to associate the events with a specific function invocation. The developer adds the following code to the Lambda function:

```
function handler(event, context) {

}
```

Which solution will meet this requirement?
A. Obtain the request identifier from the AWS request ID field in the context object. Configure the application to write logs to standard output.
B. Obtain the request identifier from the AWS request ID field in the event object. Configure the application to write logs to a file.
C. Obtain the request identifier from the AWS request ID field in the event object. Configure the application to write logs to standard output.
D. Obtain the request identifier from the AWS request ID field in the context object. Configure the application to write logs to a file.
Answer: (SHOW ANSWER)
https://docs.aws.amazon.com/lambda/latest/dg/nodejs-context.html
https://docs.aws.amazon.com/lambda/latest/dg/nodejs-logging.html There is no explicit information for the runtime, the code is written in Node.js.
AWS Lambda is a service that lets developers run code without provisioning or managing servers. The developer can use the AWS request ID field in the context object to obtain a unique identifier for each function invocation. The developer can configure the application to write logs to standard output, which will be captured by Amazon CloudWatch Logs. This solution will meet the requirement of logging key events with a unique identifier.
Reference:
[What Is AWS Lambda? - AWS Lambda]
[AWS Lambda Function Handler in Node.js - AWS Lambda]
[Using Amazon CloudWatch - AWS Lambda]

Valid DVA-C02 Dumps shared by TrainingQuiz.com for Helping Passing DVA-C02 Exam! TrainingQuiz.com now offer the **newest DVA-C02 exam dumps**, the TrainingQuiz.com DVA-C02 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C02 dumps with Test Engine here: https://www.trainingquiz.com/DVA-C02-practice-quiz.html (**649** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)