



Amazon

Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

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

Which solution will meet this requirement MOST cost-effectively?

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

Answer: C

Explanation:

Amazon S3 is a service that provides highly scalable, durable, and secure object storage. The developer can create an S3 bucket to host the repository and migrate the existing .xml files to the S3 bucket. The developer can update the application code to use the AWS SDK to read and write configuration files from S3. This solution will meet the requirement of high availability for the repository in a cost-effective way.

References:

? [Amazon Simple Storage Service (S3)]

? [Using AWS SDKs with Amazon S3]

NEW QUESTION 2

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

Which option would enable DynamoDB table updates to invoke the Lambda function?

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

Answer: B

Explanation:

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

Reference: [Using AWS Lambda with Amazon DynamoDB], [Using AWS Lambda with Amazon SNS]

NEW QUESTION 3

An application that runs on AWS receives messages from an Amazon Simple Queue Service (Amazon SQS) queue and processes the messages in batches. The application sends the data to another SQS queue to be consumed by another legacy application. The legacy system can take up to 5 minutes to process some transaction data.

A developer wants to ensure that there are no out-of-order updates in the legacy system. The developer cannot alter the behavior of the legacy system.

Which solution will meet these requirements?

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

Answer: A

Explanation:

? An SQS FIFO queue is a type of queue that preserves the order of messages and ensures that each message is delivered and processed only once1. This is suitable for the scenario where the developer wants to ensure that there are no out-of-order updates in the legacy system.

? The visibility timeout value is the amount of time that a message is invisible in the queue after a consumer receives it2. This prevents other consumers from processing the same message simultaneously. If the consumer does not delete the message before the visibility timeout expires, the message becomes visible again and another consumer can receive it2.

? In this scenario, the developer needs to configure the visibility timeout value to be longer than the maximum processing time of the legacy system, which is 5 minutes. This will ensure that the message remains invisible in the queue until the legacy system finishes processing it and deletes it. This will prevent duplicate or out-of-order processing of messages by the legacy system.

NEW QUESTION 4

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 IAM user with an appropriate policy.
- C. Store the access key ID and secret access key on the EC2 instances.
- D. Modify the application to use the S3 `GeneratePresignedUrl` API call.
- E. Modify the application to use the S3 `GetObject` API call and to return the object handle to the user.
- F. Modify the application to delegate requests to the S3 bucket.

Answer: AC

Explanation:

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

- ? Use Amazon S3 with Amazon EC2
- ? How to Access AWS S3 Bucket from EC2 Instance In a Secured Way
- ? Sharing an Object with Others

NEW QUESTION 5

A company runs a payment application on Amazon EC2 instances behind an Application Load Balance. The EC2 instances run in an Auto Scaling group across multiple Availability Zones. The application needs to retrieve application secrets during the application startup and export the secrets as environment variables. These secrets must be encrypted at rest and need to be rotated every month. Which solution will meet these requirements with the LEAST development effort?

- A. Save the secrets in a text file and store the text file in Amazon S3. Provision a customer managed key. Use the key for secret encryption in Amazon S3. Read the contents of the text file and read the export as environment variables. Configure S3 Object Lambda to rotate the text file every month.
- B. Save the secrets as strings in AWS Systems Manager Parameter Store and use the default AWS Key Management Service (AWS KMS) key. Configure an Amazon EC2 user data script to retrieve the secrets during the startup and export as environment variables. Configure an AWS Lambda function to rotate the secrets in Parameter Store every month.
- C. Save the secrets as base64 encoded environment variables in the application properties.
- D. Retrieve the secrets during the application startup.
- E. Reference the secrets in the application code.
- F. Write a script to rotate the secrets saved as environment variables.
- G. Store the secrets in AWS Secrets Manager. Provision a new customer master key. Use the key to encrypt the secrets. Enable automatic rotation. Configure an Amazon EC2 user data script to programmatically retrieve the secrets during the startup and export as environment variables.

Answer: D

Explanation:

AWS Secrets Manager is a service that enables the secure management and rotation of secrets, such as database credentials, API keys, or passwords. By using Secrets Manager, the company can avoid hardcoding secrets in the application code or properties files, and instead retrieve them programmatically during the application startup. Secrets Manager also supports automatic rotation of secrets by using AWS Lambda functions or built-in rotation templates. The company can provision a customer master key (CMK) to encrypt the secrets and use the AWS SDK or CLI to export the secrets as environment variables. References:

- ? What Is AWS Secrets Manager? - AWS Secrets Manager
- ? Rotating Your AWS Secrets Manager Secrets - AWS Secrets Manager
- ? Retrieving a Secret - AWS Secrets Manager

NEW QUESTION 6

A developer is creating an application that will store personal health information (PHI). The PHI needs to be encrypted at all times. An encrypted Amazon RDS for MySQL DB instance is storing the data. The developer wants to increase the performance of the application by caching frequently accessed data while adding the ability to sort or rank the cached datasets. Which solution will meet these requirements?

- A. Create an Amazon ElastiCache for Redis instance.
- B. Enable encryption of data in transit and at rest.
- C. Store frequently accessed data in the cache.
- D. Create an Amazon ElastiCache for Memcached instance.
- E. Enable encryption of data in transit and at rest.
- F. Store frequently accessed data in the cache.
- G. Create an Amazon RDS for MySQL read replica.
- H. Connect to the read replica by using SSL.
- I. Configure the read replica to store frequently accessed data.
- J. Create an Amazon DynamoDB table and a DynamoDB Accelerator (DAX) cluster for the table.
- K. Store frequently accessed data in the DynamoDB table.

Answer: A

Explanation:

Amazon ElastiCache is a service that offers fully managed in-memory data stores that are compatible with Redis or Memcached. The developer can create an ElastiCache for Redis instance and enable encryption of data in transit and at rest. This will ensure that the PHI is encrypted at all times. The developer can store frequently accessed data in the cache and use Redis features such as sorting and ranking to enhance the performance of the application.

References:

- ? [What Is Amazon ElastiCache? - Amazon ElastiCache]
- ? [Encryption in Transit - Amazon ElastiCache for Redis]
- ? [Encryption at Rest - Amazon ElastiCache for Redis]

NEW QUESTION 7

A developer is creating an AWS CloudFormation template to deploy Amazon EC2 instances across multiple AWS accounts. The developer must choose the EC2 instances from a list of approved instance types.

How can the developer incorporate the list of approved instance types in the CloudFormation template?

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

Answer: D

Explanation:

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

NEW QUESTION 8

A developer is working on a Python application that runs on Amazon EC2 instances. The developer wants to enable tracing of application requests to debug performance issues in the code.

Which combination of actions should the developer take to achieve this goal? (Select TWO)

- A. Install the Amazon CloudWatch agent on the EC2 instances.
- B. Install the AWS X-Ray daemon on the EC2 instances.
- C. Configure the application to write JSON-formatted logs to `/var/log/cloudwatch`.
- D. Configure the application to write trace data to `/var/log/xray`.
- E. Install and configure the AWS X-Ray SDK for Python in the application.

Answer: BE

Explanation:

This solution will meet the requirements by using AWS X-Ray to enable tracing of application requests to debug performance issues in the code. AWS X-Ray is a service that collects data about requests that the applications serve, and provides tools to view, filter, and gain insights into that data. The developer can install the AWS X-Ray daemon on the EC2 instances, which is a software that listens for traffic on UDP port 2000, gathers raw segment data, and relays it to the X-Ray API. The developer can also install and configure the AWS X-Ray SDK for Python in the application, which is a library that enables instrumenting Python code to generate and send trace data to the X-Ray daemon. Option A is not optimal because it will install the Amazon CloudWatch agent on the EC2 instances, which is a software that collects metrics and logs from EC2 instances and on-premises servers, not application performance data. Option C is not optimal because it will configure the application to write JSON-formatted logs to `/var/log/cloudwatch`, which is not a valid path or destination for CloudWatch logs. Option D is not optimal because it will configure the application to write trace data to `/var/log/xray`, which is also not a valid path or destination for X-Ray trace data.

References: [AWS X-Ray], [Running the X-Ray Daemon on Amazon EC2]

NEW QUESTION 9

A developer must use multi-factor authentication (MFA) to access data in an Amazon S3

bucket that is in another AWS account. Which AWS Security Token Service (AWS STS) API operation should the developer use with the MFA information to meet this requirement?

- A. AssumeRoleWithWebIdentity
- B. GetFederationToken
- C. AssumeRoleWithSAML
- D. AssumeRole

Answer: D

Explanation:

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

? AssumeRole

? Requesting Temporary Security Credentials

NEW QUESTION 10

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
- B. Call Amazon S3 by using a GET request to access the object without PII.
- C. Set up an S3 event notification that invokes the `removePii` function when an S3 PUT request is made

- D. Call Amazon S3 by using a PUT request to access the object without PII.
- E. Create an S3 Object Lambda access point from the S3 console
- F. Select the removePii function
- G. Use S3 Access Points to access the object without PII.
- H. Create an S3 access point from the S3 console
- I. Use the access point name to call the GetObjectLegalHold S3 API function
- J. Pass in the removePii function name to access the object without PII.

Answer: C

Explanation:

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

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

NEW QUESTION 10

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

References:

- ? [What Is AWS Secrets Manager? - AWS Secrets Manager]
- ? [Retrieving a Secret - AWS Secrets Manager]

NEW QUESTION 13

A company is migrating an on-premises database to Amazon RDS for MySQL. The company has read-heavy workloads. The company wants to refactor the code to achieve optimum read performance for queries. Which solution will meet this requirement with LEAST current and future effort?

- A. Use a multi-AZ Amazon RDS deployment
- B. Increase the number of connections that the code makes to the database or increase the connection pool size if a connection pool is in use.
- C. Use a multi-AZ Amazon RDS deployment
- D. Modify the code so that queries access the secondary RDS instance.
- E. Deploy Amazon RDS with one or more read replicas
- F. Modify the application code so that queries use the URL for the read replicas.
- G. Use open source replication software to create a copy of the MySQL database on an Amazon EC2 instance
- H. Modify the application code so that queries use the IP address of the EC2 instance.

Answer: C

Explanation:

Amazon RDS for MySQL supports read replicas, which are copies of the primary database instance that can handle read-only queries. Read replicas can improve the read performance of the database by offloading the read workload from the primary instance and distributing it across multiple replicas. To use read replicas, the application code needs to be modified to direct read queries to the URL of the read replicas, while write queries still go to the URL of the primary instance. This solution requires less current and future effort than using a multi-AZ deployment, which does not provide read scaling benefits, or using open source replication software, which requires additional configuration and maintenance. Reference: Working with read replicas

NEW QUESTION 18

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

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

Answer: A

Explanation:

This solution allows easy and secure control of access to the downloads at the lowest cost because it uses a content delivery network (CDN) that can cache and distribute firmware updates to customers around the world, and uses a mechanism that can restrict access to specific files or versions. Amazon CloudFront is a CDN that can improve performance, availability, and security of web applications by delivering content from edge locations closer to customers. Amazon S3 is a storage service that can store firmware updates in buckets and objects. Signed URLs are URLs that include additional information, such as an expiration date and time, that give users temporary access to specific objects in S3 buckets. The developer can use CloudFront to serve firmware updates from S3 buckets and use signed URLs to control who can download them and for how long. Creating a dedicated CloudFront distribution for each customer will incur unnecessary costs and

complexity. Using Amazon CloudFront with AWS Lambda@Edge will require additional programming overhead to implement custom logic at the edge locations. Using Amazon API Gateway and AWS Lambda to control access to an S3 bucket will also require additional programming overhead and may not provide optimal performance or availability.

Reference: [Serving Private Content through CloudFront], [Using CloudFront with Amazon S3]

NEW QUESTION 23

A developer is using an AWS Lambda function to generate avatars for profile pictures that are uploaded to an Amazon S3 bucket. The Lambda function is automatically invoked for profile pictures that are saved under the /original/ S3 prefix. The developer notices that some pictures cause the Lambda function to time out. The developer wants to implement a fallback mechanism by using another Lambda function that resizes the profile picture.

Which solution will meet these requirements with the LEAST development effort?

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

Answer: A

Explanation:

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

Reference: Using AWS Lambda destinations

NEW QUESTION 27

An application runs on multiple EC2 instances behind an ELB.

Where is the session data best written so that it can be served reliably across multiple requests?

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

Answer: A

Explanation:

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

Reference: Using ElastiCache for session management

NEW QUESTION 32

An application is processing clickstream data using Amazon Kinesis. The clickstream data feed into Kinesis experiences periodic spikes. The PutRecords API call occasionally fails and the logs show that the failed call returns the response shown below:

```
{
  "FailedRecordCount": 1,
  "Records": [
    {
      "SequenceNumber": "21269319989900637946712965403778482371",
      "ShardId": "shardId-000000000001"
    },
    {
      "ErrorCode": "ProvisionedThroughputExceededException",
      "ErrorMessage": "Rate exceeded for shard shardId-000000000001 in
        stream exampleStreamName under account 123456789."
    },
    {
      "SequenceNumber": "21269319989999637946712965403778482985",
      "ShardId": "shardId-000000000002"
    }
  ]
}
```

Which techniques will help mitigate this exception? (Choose two.)

- A. Implement retries with exponential backoff.
- B. Use a PutRecord API instead of PutRecords.

- C. Reduce the frequency and/or size of the requests.
- D. Use Amazon SNS instead of Kinesis.
- E. Reduce the number of KCL consumers.

Answer: AC

Explanation:

The response from the API call indicates that the ProvisionedThroughputExceededException exception has occurred. This exception means that the rate of incoming requests exceeds the throughput limit for one or more shards in a stream. To mitigate this exception, the developer can use one or more of the following techniques:

- ? Implement retries with exponential backoff. This will introduce randomness in the retry intervals and avoid overwhelming the shards with retries.
- ? Reduce the frequency and/or size of the requests. This will reduce the load on the shards and avoid throttling errors.
- ? Increase the number of shards in the stream. This will increase the throughput capacity of the stream and accommodate higher request rates.
- ? Use a PutRecord API instead of PutRecords. This will reduce the number of records per request and avoid exceeding the payload limit.

References:

- ? [ProvisionedThroughputExceededException - Amazon Kinesis Data Streams Service API Reference]
- ? [Best Practices for Handling Kinesis Data Streams Errors]

NEW QUESTION 33

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

- A. Create an Amazon Kinesis data stream, and attach it to the DynamoDB table.
- B. Create a trigger to connect the data stream to the Lambda function.
- C. Create an Amazon EventBridge rule to invoke the Lambda function on a regular schedule.
- D. Connect the DynamoDB table to the Lambda function to detect changes.
- E. Enable DynamoDB Streams on the table.
- F. Create a trigger to connect the DynamoDB stream to the Lambda function.
- G. Create an Amazon Kinesis Data Firehose delivery stream, and attach it to the DynamoDB table.
- H. Configure the delivery stream destination as the Lambda function.

Answer: C

Explanation:

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

References:

- ? [Amazon DynamoDB]
- ? [DynamoDB Streams - Amazon DynamoDB]
- ? [Using AWS Lambda with Amazon DynamoDB - AWS Lambda]

NEW QUESTION 36

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 to 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: BD

Explanation:

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.

References: [Using HTTPS with CloudFront], [Restricting Access to Amazon S3 Content by Using an Origin Access Identity]

NEW QUESTION 38

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: B

Explanation:

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 39

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 to an Amazon DynamoDB database.
- C. Configure the Amazon Aurora MySQL DB instance for Multi-AZ deployment.
- D. Create a proxy in Amazon RDS Proxy Query the proxy instead of the DB instance.

Answer: D

Explanation:

This solution will meet the requirements by using Amazon RDS Proxy, which is a fully managed, highly available database proxy for Amazon RDS that makes applications more scalable, more resilient to database failures, and more secure. The developer can create a proxy in Amazon RDS Proxy, which sits between the application

and the DB instance and handles connection management, pooling, and routing. The developer can query the proxy instead of the DB instance, which reduces the number of open connections to the DB instance and avoids errors for too many connections. Option A is not optimal because it will create a read replica for the DB instance, which may not solve the problem of too many connections as read replicas also have connection limits and may incur additional costs. Option B is not optimal because it will migrate the data to an Amazon DynamoDB database, which may introduce additional complexity and overhead for migrating and accessing data from a different database service. Option C is not optimal because it will configure the Amazon Aurora MySQL DB instance for Multi-AZ deployment, which may improve availability and durability of the DB instance but not reduce the number of connections.

References: [Amazon RDS Proxy], [Working with Amazon RDS Proxy]

NEW QUESTION 43

A company's developer has deployed an application in AWS by using AWS CloudFormation The CloudFormation stack includes parameters in AWS Systems Manager Parameter Store that the application uses as configuration settings. The application can modify the parameter values

When the developer updated the stack to create additional resources with tags, the developer noted that the parameter values were reset and that the values ignored the latest changes made by the application. The developer needs to change the way the company deploys the CloudFormation stack. The developer also needs to avoid resetting the parameter values outside the stack.

Which solution will meet these requirements with the LEAST development effort?

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

Answer: D

Explanation:

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

NEW QUESTION 46

A company needs to set up secure database credentials for all its AWS Cloud resources. The company's resources include Amazon RDS DB instances Amazon DocumentDB clusters and Amazon Aurora DB instances. The company's security policy mandates that database credentials be encrypted at rest and rotated at a regular interval.

Which solution will meet these requirements MOST securely?

- A. Set up IAM database authentication for token-based access
- B. Generate user tokens to provide centralized access to RDS DB instance
- C. Amazon DocumentDB clusters and Aurora DB instances.
- D. Create parameters for the database credentials in AWS Systems Manager Parameter Store Set the Type parameter to Secure String
- E. Set up automatic rotation on the parameters.
- F. Store the database access credentials as an encrypted Amazon S3 object in an S3 bucket Block all public access on the S3 bucket automatic rotation on the encryption key.
- G. Use S3 server-side encryption to set up
- H. Create an AWS Lambda function by using the SecretsManagerRotationTemplate template in the AWS Secrets Manager console
- I. Create secrets for the database credentials in Secrets Manager Set up secrets rotation on a schedule.

Answer: D

Explanation:

This solution will meet the requirements by using AWS Secrets Manager, which is a service that helps protect secrets such as database credentials by encrypting them with AWS Key Management Service (AWS KMS) and enabling automatic rotation of secrets. The developer can create an AWS Lambda function by using the SecretsManagerRotationTemplate template in the AWS Secrets Manager console, which provides a sample code for rotating secrets for RDS DB instances, Amazon DocumentDB clusters, and Amazon Aurora DB instances. The developer can also create secrets for the database credentials in Secrets Manager, which encrypts them at rest and provides secure access to them. The developer can set up secrets rotation on a schedule, which changes the database credentials periodically according to a specified interval or event. Option A is not optimal because it will set up IAM database authentication for token-based access, which may not be compatible with all database engines and may require additional configuration and management of IAM roles or users. Option B is not optimal because it will create parameters for the database credentials in AWS Systems Manager Parameter Store, which does not support automatic rotation of secrets. Option C is not optimal because it will store the database access credentials as an encrypted Amazon S3 object in an S3 bucket, which may introduce additional costs and complexity for accessing and securing the data.

References: [AWS Secrets Manager], [Rotating Your AWS Secrets Manager Secrets]

NEW QUESTION 51

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-digit universally unique identifier (UUID)
- A Product ID and User ID 16 digit UUIDs that reference other tables
- A Product Rating on a scale of 1-5
- An optional comment from the user

The table partition key is the Review ID. The most performed query against the table is to find the 10 reviews with the highest rating for a given product. Which index will provide the FASTEST response for this query?"

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

Answer: A

Explanation:

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

Reference: [Global Secondary Indexes], [Querying]

NEW QUESTION 56

A developer is working on a web application that uses Amazon DynamoDB as its data store. The application has two DynamoDB tables: one table named artists and one table named songs. The artists table has artistName as the partition key. The songs table has songName as the partition key and artistName as the sort key.

The table usage patterns include the retrieval of multiple songs and artists in a single database operation from the webpage. The developer needs a way to retrieve this information with minimal network traffic and optimal application performance.

Which solution will meet these requirements?"

- A. Perform a BatchGetItem operation that returns items from the two tables.
- B. Use the list of songName/artistName keys for the songs table and the list of artistName key for the artists table.
- C. Create a local secondary index (LSI) on the songs table that uses artistName as the partition key. Perform a query operation for each artistName on the songs table that filters by the list of songName. Perform a query operation for each artistName on the artists table.
- D. Perform a BatchGetItem operation on the songs table that uses the songName/artistName key.
- E. Perform a BatchGetItem operation on the artists table that uses artistName as the key.
- F. Perform a Scan operation on each table that filters by the list of songName/artistName for the songs table and the list of artistName in the artists table.

Answer: A

Explanation:

BatchGetItem can return one or multiple items from one or more tables. For reference, check the link below:
https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_BatchGetItem.html

NEW QUESTION 57

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

Answer: A

Explanation:

The most cost-effective solution is to use the DefinitionSubstitutions property of the AWS::StepFunctions::StateMachine resource to inject the API endpoint as a variable in the state machine definition. This way, the developer can use the intrinsic function

Fn::GetAtt to get the API endpoint from the AWS::ApiGateway::RestApi resource, and pass it to the state machine without creating any additional resources or environment variables. The other solutions involve creating and managing extra resources, such as Secrets Manager secrets or AppConfig configuration profiles, which incur additional costs and complexity. References

? AWS::StepFunctions::StateMachine - AWS CloudFormation

? Call API Gateway with Step Functions - AWS Step Functions

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

NEW QUESTION 61

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

Answer: A

Explanation:

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

NEW QUESTION 65

A company is running a custom application on a set of on-premises Linux servers that are accessed using Amazon API Gateway. AWS X-Ray tracing has been enabled on the API test stage.

How can a developer enable X-Ray tracing on the on-premises servers with the LEAST amount of configuration?

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

Answer: B

Explanation:

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

References:

? [AWS X-Ray concepts - AWS X-Ray]

? [Setting up AWS X-Ray - AWS X-Ray]

NEW QUESTION 69

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
- C. Configure the agent to use IAM user credentials with permissions for CloudWatch.
- D. Upload log files from the on-premises server to Amazon S3 and have CloudWatch read the files.
- E. Upload log files from the on-premises server to an Amazon EC2 instance and have the instance forward the logs to CloudWatch.

Answer: B

Explanation:

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.

References:

? [What Is Amazon CloudWatch? - Amazon CloudWatch]

? [Installing and Configuring the CloudWatch Agent - Amazon CloudWatch]

NEW QUESTION 71

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?

Query the instance metadata from `http://169.254.169.254/latest/meta-data/`.

- A. 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

Explanation:

The instance metadata service provides information about the EC2 instance, including the public IPv4 address, which can be obtained by querying the endpoint `http://169.254.169.254/latest/meta-data/public-ipv4`.
References
? Instance metadata and user data
? Get Public IP Address on current EC2 Instance
? Get the public ip address of your EC2 instance quickly

NEW QUESTION 73

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

- A. Store the third-party service endpoints in Lambda layers that correspond to the stage
- B. Store the third-party service endpoints in API Gateway stage variables that correspond to the stage
- C. Encode the third-party service endpoints as query parameters in the API Gateway request URL.
- D. Store the third-party service endpoint for each environment in AWS AppConfig

Answer: B

Explanation:

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

NEW QUESTION 77

A developer wants to add request validation to a production environment Amazon API Gateway API. The developer needs to test the changes before the API is deployed to the production environment. For the least 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
- B. Create a new API Import the OpenAPI file Modify the new API to add request validation
- C. Perform the tests Modify the existing API to add request validation
- D. Deploy the existing API to production.
- E. Modify the existing API to add request validation
- F. Deploy the updated API to a new API Gateway stage Perform the tests Deploy the updated API to the API Gateway production stage.
- G. Create a new API Add the necessary resources and methods including new request validation
- H. Perform the tests Modify the existing API to add request validation
- I. Deploy the existing API to production.
- J. Clone the existing API Modify the new API to add request validation
Modify the existing API to add request validation Deploy the existing API to production.
- K. Perform the tests

Answer: D

Explanation:

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

NEW QUESTION 79

A developer migrated a legacy application to an AWS Lambda function. The function uses a third-party service to pull data with a series of API calls at the end of each month. The function then processes the data to generate the monthly reports. The function has been working with no issues so far. The third-party service recently issued a restriction to allow a feed number to API calls each minute and each day. If the API calls exceed the limit for each minute or each day, then the service will produce errors. The API also provides the minute limit and daily limit in the response header. This restriction might extend the overall process to multiple days because the process is consuming more API calls than the available limit. What is the MOST operationally efficient way to refactor the serverless application to accommodate this change?

- A. Use an AWS Step Functions State machine to monitor API failure
- B. Use the Wait state to delay calling the Lambda function.
- C. Use an Amazon Simple Queue Service (Amazon SQS) queue to hold the API call
- D. Configure the Lambda function to poll the queue within the API threshold limits.
Use an Amazon CloudWatch Logs metric to count the number of API call
- E. Configure an Amazon CloudWatch alarm that stops the currently running instance of the Lambda function when the metric exceeds the API threshold limits.
- G. Use Amazon Kinesis Data Firehose to batch the API calls and deliver them to an Amazon S3 bucket with an event notification to invoke the Lambda function.

Answer: A

Explanation:

The solution that will meet the requirements is to use an AWS Step Functions state machine to monitor API failures. Use the Wait state to delay calling the Lambda function. This way, the developer can refactor the serverless application to accommodate the change in a way that is automated and scalable. The developer can use Step Functions to orchestrate the Lambda function and handle any errors or retries. The developer can also use the Wait state to pause the execution for a specified duration or until a specified timestamp, which can help avoid exceeding the API limits. The other options either involve using additional services that are not necessary or appropriate for this scenario, or do not address the issue of API failures.

Reference: AWS Step Functions Wait state

NEW QUESTION 82

A developer is building a web application that uses Amazon API Gateway to expose an AWS Lambda function to process requests from clients. During testing, the developer notices that the API Gateway times out even though the Lambda function finishes under the set time limit.

Which of the following API Gateway metrics in Amazon CloudWatch can help the developer troubleshoot the issue? (Choose two.)

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

Answer: BD

Explanation:

Amazon API Gateway is a service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. Amazon CloudWatch is a service that monitors AWS resources and applications. API Gateway provides several CloudWatch metrics to help developers troubleshoot issues with their APIs. Two of the metrics that can help the developer troubleshoot the issue of API Gateway timing out are:

? IntegrationLatency: This metric measures the time between when API Gateway relays a request to the backend and when it receives a response from the backend. A high value for this metric indicates that the backend is taking too long to respond and may cause API Gateway to time out.

? Latency: This metric measures the time between when API Gateway receives a request from a client and when it returns a response to the client. A high value for this metric indicates that either the integration latency is high or API Gateway is taking too long to process the request or response.

References:

? [What Is Amazon API Gateway? - Amazon API Gateway]

? [Amazon API Gateway Metrics and Dimensions - Amazon CloudWatch]

? [Troubleshooting API Errors - Amazon API Gateway]

NEW QUESTION 87

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

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

Answer: C

Explanation:

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

Reference: How Application Load Balancer works with your applications

NEW QUESTION 90

A developer is creating a mobile app that calls a backend service by using an Amazon API Gateway REST API. For integration testing during the development phase, the developer wants to simulate different backend responses without invoking the backend service.

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 93

A development team wants to build a continuous integration/continuous delivery (CI/CD) pipeline. The team is using AWS CodePipeline to automate the code build and deployment. The team wants to store the program code to prepare for the CI/CD pipeline. Which AWS service should the team use to store the program code?

- A. AWS CodeDeploy
- B. AWS CodeArtifact
- C. AWS CodeCommit
- D. Amazon CodeGuru

Answer: C

Explanation:

AWS CodeCommit is a service that provides fully managed source control for hosting secure and scalable private Git repositories. The development team can use CodeCommit to store the program code and prepare for the CI/CD pipeline. CodeCommit integrates with other AWS services such as CodePipeline, CodeBuild, and CodeDeploy to automate the code build and deployment process.

References:

? [What Is AWS CodeCommit? - AWS CodeCommit]

? [AWS CodePipeline - AWS CodeCommit]

NEW QUESTION 98

A company is using Amazon OpenSearch Service to implement an audit monitoring system. A developer needs to create an AWS CloudFormation custom resource that is

associated with an AWS Lambda function to configure the OpenSearch Service domain. The Lambda function must access the OpenSearch Service domain by using Open Search Service internal master user credentials. What is the MOST secure way to pass these credentials to the Lambda function?

- A. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and the Lambda function's environment variable
- B. Set the No Echo attribute to true.
- C. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and to create a parameter
- D. In AWS Systems Manager Parameter Store
- E. Set the No Echo attribute to true
- F. Create an IAM role that has the ssm:GetParameter permission
- G. Assign the role to the Lambda function
- H. Store the parameter name as the Lambda function's environment variable
- I. Resolve the parameter's value at runtime.
- J. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and the Lambda function's environment variable. We encrypt the parameter's value by using the AWS Key Management Service (AWS KMS) Encrypt command.
- K. Use CloudFormation to create an AWS Secrets Manager secret
- L. Use a CloudFormation dynamic reference to retrieve the secret's value for the OpenSearch Service domain's MasterUserOptions
- M. Create an IAM role that has the secrets:manage
- N. GetSecretValue permission
- O. Assign the role to the Lambda Function. Store the secret's name as the Lambda function's environment variable
- P. Resolve the secret's value at runtime.

Answer: D

Explanation:

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

Reference: Using dynamic references to specify template values

NEW QUESTION 103

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.
- C. Create API Gateway resources and set the integration type value to AWS_PROXY. Deploy the API.
- D. Create an EC2 application that returns mocked HTTP responses. Create API Gateway resources and set the integration type value to AWS. Create an API Gateway stage and deploy the API.
- E. Create API Gateway resources and set the integration type value set to HTTP_PROXY.
- F. Add mapping templates and deploy the API.
- G. Create an AWS Lambda layer that returns various HTTP status codes. Associate the Lambda layer with the API deployment.

Answer: A

Explanation:

The best solution for publishing an API without an integrated backend is to use the MOCK integration type in API Gateway. This allows the developer to return a static response to the client without sending the request to a backend service. The developer can configure the method integration request and integration response to associate a response with an HTTP status code, such as 200 OK or 404 Not Found. The developer can also create an API Gateway stage and deploy

the API to make it available to the teams that depend on the application backend. The other solutions are either not feasible or not efficient. Creating an AWS Lambda function, an EC2 application, or an AWS Lambda layer would require additional resources and code to generate the mocked responses and HTTP status codes. These solutions would also incur additional costs and complexity, and would not leverage the built-in functionality of API Gateway. References

- ? Set up mock integrations for API Gateway REST APIs
- ? Mock Integration for API Gateway - AWS CloudFormation
- ? Mocking API Responses with API Gateway
- ? How to mock API Gateway responses with AWS SAM

NEW QUESTION 106

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 size or complexity of the files. The UI team wants the API to provide an immediate response so that the UI can display 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 with a static 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.
- C. Configure the maximum age of the event so that the Lambda function will run asynchronously.
- D. Change the API Gateway timeout value to match the Lambda function timeout value.
- E. Deploy the API Gateway stage to apply the changes.
- F. Change the API Gateway route to add an X-Amz-Target header with a static value of 'A sync' in the integration request. Deploy the API Gateway stage to apply the changes.

Answer: A

Explanation:

This solution allows the API to invoke the Lambda function asynchronously, which means that the API will return an immediate response without waiting for the function to complete. The X-Amz-Invocation-Type header specifies the invocation type of the Lambda function, and setting it to 'Event' means that the function will be invoked asynchronously. The function can then use Amazon Simple Email Service (SES) to send an email message when the report processing is complete.

Reference: [Asynchronous invocation], [Set up Lambda proxy integrations in API Gateway]

NEW QUESTION 109

A developer is designing a serverless application for a game in which users register and log in through a web browser. The application makes requests on behalf of users to a set of AWS Lambda functions that run behind an Amazon API Gateway HTTP API.

The developer needs to implement a solution to register and log in users on the application's sign-in page. The solution must minimize operational overhead and must minimize ongoing management of user identities.

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 114

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 in Amazon S3 and use S3 Transfer Acceleration.
- C. Add retries with exponential back off for DynamoDB queries.
- D. Use DynamoDB Accelerator (DAX) to cache the trading data.

Answer: D

Explanation:

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

References: [DynamoDB Accelerator (DAX)], [Local Secondary Indexes]

NEW QUESTION 117

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 the 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 in AWS Systems Manager Parameter Store. Modify the Lambda function to take the latest request response from Parameter Store.

Answer: A

Explanation:

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

References: [Amazon CloudFront], [Caching Content Based on Request Headers]

NEW QUESTION 121

A developer wants to store information about movies. Each movie has a title, release year, and genre. The movie information also can include additional properties about the cast and production crew. This additional information is inconsistent across movies. For example, one movie might have an assistant director, and another movie might have an animal trainer.

The developer needs to implement a solution to support the following use cases:

For a given title and release year, get all details about the movie that has that title and release year.

For a given title, get all details about all movies that have that title. For a given genre, get all details about all movies in that genre. Which data store configuration will meet these requirements?

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

Answer: A

Explanation:

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

References:

? [Amazon DynamoDB]

? [Working with Queries - Amazon DynamoDB]

? [Working with Global Secondary Indexes - Amazon DynamoDB]

NEW QUESTION 123

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

Which solution will meet these requirements'?

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

? Amazon ElastiCache for Redis

? Caching Strategies and Best Practices - Amazon ElastiCache for Redis

? Using Amazon ElastiCache for Redis with Amazon RDS

? Amazon DynamoDB Accelerator (DAX)

? Amazon S3 Transfer Acceleration

? Amazon CloudFront

NEW QUESTION 125

.....

Relate Links

100% Pass Your AWS-Certified-Developer-Associate Exam with Exambible Prep Materials

<https://www.exambible.com/AWS-Certified-Developer-Associate-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>