

Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate

<https://www.2passeasy.com/dumps/AWS-Certified-Developer-Associate/>



NEW QUESTION 1

- (Exam Topic 1)

A Developer needs temporary access to resources in a second account. What is the MOST secure way to achieve this?

- A. Use the Amazon Cognito user pools to get short-lived credentials for the second account.
- B. Create a dedicated IAM access key for the second account, and send it by mail.
- C. Create a cross-account access role, and use sts:AssumeRole API to get short-lived credentials.
- D. Establish trust, and add an SSH key for the second account to the IAM user.

Answer: C

NEW QUESTION 2

- (Exam Topic 1)

A Developer has been asked to make changes to the source code of an AWS Lambda function. The function is managed using an AWS CloudFormation template. The template is configured to load the source code from an Amazon S3 bucket. The Developer manually created a .ZIP file deployment package containing the changes and put the file into the correct location on Amazon S3. When the function is invoked, the code changes have not been applied. What step is required to update the function with the changes?

- A. Delete the .ZIP file on S3, and re-upload by using a different object key name.
- B. Update the CloudFormation stack with the correct values for the function code properties S3Bucket, S3Key, or S3ObjectVersion.
- C. Ensure that the function source code is base64-encoded before uploading the deployment package to S3.
- D. Modify the execution role of the Lambda function to allow S3 access permission to the deployment package .ZIP file.

Answer: B

Explanation:

Changes to a deployment package in Amazon S3 are not detected automatically during stack updates. To update the function code, change the object key or version in the template.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-lambda-function-code.htm>

NEW QUESTION 3

- (Exam Topic 1)

An e-commerce site allows returning users to log in to display customized web pages. The workflow is shown in the image below:



An application is running on EC2 instances. Amazon RDS is used for the database that stores user accounts and preferences. The website freezes or is slow to load while waiting for the login step to complete. The remaining components of the site are well-optimized.

Which of the following techniques will resolve this issue? (Select Two.)

- A. Implement the user login page as an asynchronous Lambda function.
- B. Use Amazon ElastiCache for MemCached to cache user data.
- C. Use Amazon Application Load Balancer to load balance the traffic to the website.
- D. Call the database asynchronously so the code can continue executing.
- E. Batch login requests from hundreds of users together as a single read request to the database.

Answer: BD

Explanation:

<https://docs.aws.amazon.com/sdk-for-javascript/v2/developer-guide/making-asynchronous-calls.html>

NEW QUESTION 4

- (Exam Topic 1)

A Developer is writing a serverless application that requires that an AWS Lambda function be invoked every 10 minutes. What is an automated and serverless way to trigger the function?

- A. Deploy an Amazon EC2 instance based on Linux, and edit its `/etc/crontab` file by adding a command to periodically invoke the Lambda function.
- B. Configure an environment variable named `PERIOD` for the Lambda function.
- C. Set the value to 600.
- D. Create an Amazon CloudWatch Events rule that triggers on a regular schedule to invoke the Lambda function.
- E. Create an Amazon SNS topic that has a subscription to the Lambda function with a 600-second timer.

Answer: C

Explanation:

Reference:

[https://aws.amazon.com/blogs/architecture/a-serverless-solution-for-invoking-aws-lambda-at-a-sub-minute-frequ](https://aws.amazon.com/blogs/architecture/a-serverless-solution-for-invoking-aws-lambda-at-a-sub-minute-frequency/)

NEW QUESTION 5

- (Exam Topic 1)

A Developer must re-implement the business logic for an order fulfillment system. The business logic has to make requests to multiple vendors to decide where to purchase an item. The whole process can take up to a week to complete.

What is the MOST efficient and SIMPLEST way to implement a system that meets these requirements?

- A. Use AWS Step Functions to execute parallel Lambda functions, and join the results.
- B. Create an AWS SQS for each vendor, poll the queue from a worker instance, and join the results.
- C. Use AWS Lambda to asynchronously call a Lambda function for each vendor, and join the results.
- D. Use Amazon CloudWatch Events to orchestrate the Lambda functions.

Answer: A

Explanation:

<https://aws.amazon.com/step-functions/>

NEW QUESTION 6

- (Exam Topic 1)

A Developer must deploy a new AWS Lambda function using an AWS CloudFormation template. Which procedures will deploy a Lambda function? (Select TWO.)

- A. Upload the code to an AWS CodeCommit repository, then add a reference to it in an `AWS::Lambda::Function` resource in the template.
- B. Create an `AWS::Lambda::Function` resource in the template, then write the code directly inside the CloudFormation template.
- C. Upload a .ZIP file containing the function code to Amazon S3, then add a reference to it in an `AWS::Lambda::Function` resource in the template.
- D. Upload a .ZIP file to AWS CloudFormation containing the function code, then add a reference to it in an `AWS::Lambda::Function` resource in the template.
- E. Upload the function code to a private Git repository, then add a reference to it in an `AWS::Lambda::Function` resource in the template.

Answer: BC

Explanation:

<https://aws.amazon.com/blogs/infrastructure-and-automation/deploying-aws-lambda-functions-using-aws-cloudf>

NEW QUESTION 7

- (Exam Topic 1)

A startup's photo-sharing site is deployed in a VPC. An ELB distributes web traffic across two subnets. ELB session stickiness is configured to use the AWS-generated session cookie, with a session TTL of 5 minutes. The webserver Auto Scaling Group is configured as: `min-size=4`, `max-size=4`.

The startup is preparing for a public launch, by running load-testing software installed on a single EC2 instance running in `us-west-2a`. After 60 minutes of load-testing, the webserver logs show:

Which recommendations can help ensure load-testing HTTP requests are evenly distributed across the four web servers? Choose 2 answers

- A. Launch and run the load-tester EC2 instance from `us-east-1` instead.
- B. Re-configure the load-testing software to re-resolve DNS for each web request.
- C. Use a 3rd-party load-testing service which offers globally-distributed test clients.
- D. Configure ELB and Auto Scaling to distribute across `us-west-2a` and `us-west-2c`.
- E. Configure ELB session stickiness to use the app-specific session cookie.

Answer: CE

NEW QUESTION 8

- (Exam Topic 1)

How can software determine the public and private IP addresses of the Amazon EC2 instance that it is running on?

- A. Query the appropriate Amazon CloudWatch metric.
- B. Use `ipconfig` or `ifconfig` command.
- C. Query the local instance user data.
- D. Query the local instance metadata.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-categories.html>

Data: `local-hostname` Descriptions: The private IPv4 DNS hostname of the instance. In cases where multiple network interfaces are present, this refers to the `eth0` device (the device for which the device number is 0). Data: `public-ipv4` Descriptions: The public IPv4 address. If an Elastic IP address is associated with the instance, the value returned is the Elastic IP address.

NEW QUESTION 9

- (Exam Topic 1)

A company is using Amazon API Gateway to manage access to a set of microservices implemented as AWS Lambda functions. Following a bug report, the company makes a minor breaking change to one of the APIs. In order to avoid impacting existing clients when the new API is deployed, the company wants to allow clients six months to migrate from v1 to v2.

Which approach should the Developer use to handle this change?

- A. Update the underlying Lambda function and provide clients with the new Lambda invocation URL.
- B. Use API Gateway to automatically propagate the change to clients, specifying 180 days in the phased deployment parameter.
- C. Use API Gateway to deploy a new stage named v2 to the API and provide users with its URL.
- D. Update the underlying Lambda function, create an Amazon CloudFront distribution with the updated Lambda function as its origin.

Answer: C

NEW QUESTION 10

- (Exam Topic 1)

Which of the following statements about SWF are true? Choose 3 answers

- A. SWF tasks are assigned once and never duplicated
- B. SWF requires an S3 bucket for workflow storage
- C. SWF workflow executions can last up to a year
- D. SWF triggers SNS notifications on task assignment
- E. SWF uses deciders and workers to complete tasks
- F. SWF requires at least 1 EC2 instance per domain

Answer: ACE

NEW QUESTION 10

- (Exam Topic 1)

If a message is retrieved from a queue in Amazon SQS, how long is the message inaccessible to other users by default?

- A. 0 seconds
- B. 1 hour
- C. 1 day
- D. forever
- E. 30 seconds

Answer: E

Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html> Visibility timeout:
default value = 30 seconds, minimum = 0 seconds, maximum = 12 hours

NEW QUESTION 12

- (Exam Topic 1)

company needs a fully-managed source control service that will work in AWS. The service must ensure that revision control synchronizes multiple distributed repositories by exchanging sets of changes peer-to-peer. All users need to work productively even when not connected to a network.

Which source control service should be used?

- A. Subversion
- B. AWS CodeBuild
- C. AWS CodeCommit
- D. AWS CodeStar

Answer: C

NEW QUESTION 17

- (Exam Topic 1)

A Developer has published an update to an application that is served to a global user base using Amazon CloudFront. After deploying the application, users are not able to see the updated changes.

How can the Developer resolve this issue?

- A. Remove the origin from the CloudFront configuration and add it again.
- B. Disable forwarding of query strings and request headers from the CloudFront distribution configuration.
- C. Invalidate all the application objects from the edge caches.
- D. Disable the CloudFront distribution and enable it again to update all the edge locations.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Invalidation.html>

NEW QUESTION 18

- (Exam Topic 1)

A Developer has created a Lambda function and is finding that the function is taking longer to complete than expected. After some debugging, the Developer has discovered that increasing compute capacity would improve performance.

How can the Developer increase the Lambda compute resources?

- A. Run on a larger instance size with more compute capacity.

- B. Increase the maximum execution time.
- C. Specify a larger compute capacity when calling the Lambda function.
- D. Increase the allocated memory for the Lambda function.

Answer: D

NEW QUESTION 20

- (Exam Topic 1)

A company wants to implement a continuous integration for its workloads on AWS. The company wants to trigger unit test in its pipeline for commits on its code repository, and wants to be notified of failure events in the pipeline. How can these requirements be met?

- A. Store the source code in AWS CodeComm
- B. Create a CodePipeline to automate unit testin
- C. Use Amazon SNS to trigger notifications of failure events.
- D. Store the source code in GitHu
- E. Create a CodePipeline to automate unit testin
- F. Use Amazon SES to trigger notifications of failure events.
- G. Store the source code on GitHu
- H. Create a CodePipeline to automate unit testin
- I. Use Amazon CloudWatch to trigger notifications of failure events.
- J. Store the source code in AWS CodeComm
- K. Create a CodePipeline to automate unit testin
- L. Use Amazon CloudWatch to trigger notification of failure events.

Answer: D

NEW QUESTION 22

- (Exam Topic 1)

During non-peak hours, a Developer wants to minimize the execution time of a full Amazon DynamoDB table scan without affecting normal workloads. The workloads average half of the strongly consistent read capacity units during non-peak hours. How would the Developer optimize this scan?

- A. Use parallel scans while limiting the rate
- B. Use sequential scans
- C. Increase read capacity units during the scan operation
- D. Change consistency to eventually consistent during the scan operation

Answer: A

Explanation:

<https://aws.amazon.com/blogs/developer/rate-limited-scans-in-amazon-dynamodb/>

NEW QUESTION 27

- (Exam Topic 1)

A Developer is creating an application that needs to locate the public IPv4 address of the Amazon EC2 instance on which it runs. How can the application locate this information?

- A. Get the instance metadata by retrieving <http://169.254.169.254/latest/metadata/>.
- B. Get the instance user data by retrieving <http://169.254.169.254/latest/userdata/>.
- C. Get the application to run IFCONFIG to get the public IP address.
- D. Get the application to run IPCONFIG to get the public IP address.

Answer: A

NEW QUESTION 30

- (Exam Topic 1)

A legacy service has an XML-based SOAP interface. The Developer wants to expose the functionality of the service to external clients with the Amazon API Gateway. Which technique will accomplish this?

- A. Create a RESTful API with the API Gateway; transform the incoming JSON into a valid XML message for the SOAP interface using mapping templates.
- B. Create a RESTful API with the API Gateway; pass the incoming JSON to the SOAP interface through an Application Load Balancer.
- C. Create a RESTful API with the API Gateway; pass the incoming XML to the SOAP interface through an Application Load Balancer.
- D. Create a RESTful API with the API Gateway; transform the incoming XML into a valid message for the SOAP interface using mapping templates.

Answer: A

Explanation:

<https://blog.codecentric.de/en/2016/12/serverless-soap-legacy-api-integration-java-aws-lambda-aws-api-gateway>

NEW QUESTION 34

- (Exam Topic 1)

When writing a Lambda function, what is the benefit of instantiating AWS clients outside the scope of the handler?

- A. Legibility and stylistic convention
- B. Taking advantage of connection re-use
- C. Better error handling
- D. Creating a new instance per invocation

Answer: B

NEW QUESTION 39

- (Exam Topic 1)

Which features can be used to restrict access to data in S3? Choose 2 answers

- A. Use S3 Virtual Hosting
- B. Set an S3 Bucket policy.
- C. Enable IAM Identity Federation.
- D. Set an S3 ACL on the bucket or the object.
- E. Create a CloudFront distribution for the bucket

Answer: BD

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/secure-s3-resources/>

NEW QUESTION 43

- (Exam Topic 1)

A Developer is developing an application that manages financial transactions. To improve security, multi-factor authentication (MFA) will be required as part of the login protocol.

What services can the Developer use to meet these requirements?

- A. Amazon DynamoDB to store MFA session data, and Amazon SNS to send MFA codes
- B. Amazon Cognito with MFA
- C. AWS Directory Service
- D. AWS IAM with MFA enabled

Answer: B

Explanation:

AWS documentation - Cognito MFA Managing Security

You can add multi-factor authentication (MFA) to a user pool to protect the identity of your users. MFA adds a second authentication method that doesn't rely solely on user name and password. You can choose to use SMS text messages, or time-based one-time (TOTP) passwords as second factors in signing in your users. You can also use adaptive authentication with its risk-based model to predict when you might need another authentication factor. It's part of the user pool advanced security features, which also include protections against compromised credentials.

NEW QUESTION 46

- (Exam Topic 1)

A Developer needs to design an application running on AWS that will be used to consume Amazon SQS messages that range from 1 KB up to 1GB in size.

How should the Amazon SQS messages be managed?

- A. Use Amazon S3 and the Amazon SQS CLI.
- B. Use Amazon S3 and the Amazon SQS Extended Client Library for Java.
- C. Use Amazon EBS and the Amazon SQS CLI.
- D. Use Amazon EFS and the Amazon SQS CLI.

Answer: B

Explanation:

Reference: <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqslimits.html>

NEW QUESTION 49

- (Exam Topic 1)

You have written an application that uses the Elastic Load Balancing service to spread traffic to several web servers. Your users complain that they are sometimes forced to login again in the middle of using your application, after they have already logged in. This is not behavior you have designed.

What is a possible solution to prevent this happening?

- A. Use instance memory to save session state.
- B. Use instance storage to save session state.
- C. Use EBS to save session state
- D. Use ElastiCache to save session state.
- E. Use Glacier to save session slate.

Answer: D

Explanation:

<https://aws.amazon.com/caching/session-management/>

NEW QUESTION 51

- (Exam Topic 1)

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 SQS to save session data.
- C. Use Amazon DynamoDB to perform scalable session hadling.

D. Use Elastic Load Balancer connection draining to stop sending requests to failing instances.

Answer: C

NEW QUESTION 52

- (Exam Topic 1)

An application is real-time processing millions of events that are received through an API.

What service could be used to allow multiple consumers to process the data concurrently and MOST cost-effectively?

- A. Amazon SNS with fanout to an SQS queue for each application
- B. Amazon SNS with fanout to an SQS FIFO (first-in, first-out) queue for each application
- C. Amazon Kinesis Firehouse
- D. Amazon Kinesis Streams

Answer: D

NEW QUESTION 53

- (Exam Topic 1)

An Amazon S3 bucket, "myawsbucket" is configured with website hosting in Tokyo region, what is the region-specific website endpoint?

- A. www.myawsbucket.ap-northeast-1.amazonaws.com
- B. myawsbucket.s3-website-ap-northeast-1.amazonaws.com
- C. myawsbucket.amazonaws.com
- D. myawsbucket.tokyo.amazonaws.com

Answer: B

Explanation:

Depending on your Region, your Amazon S3 website endpoint follows one of these two formats. s3-website dash (-) Region <http://bucket-name.s3-website-Region.amazonaws.com>

s3-website dot (.) Region

<http://bucket-name.s3-website.Region.amazonaws.com>

<https://docs.aws.amazon.com/AmazonS3/latest/dev/WebsiteEndpoints.html>

NEW QUESTION 57

- (Exam Topic 1)

A game stores user game data in an Amazon DynamoDB table. Individual users should not have access to other users' game data. How can this be accomplished?

- A. Encrypt the game data with individual user keys.
- B. Restrict access to specific items based on certain primary key values.
- C. Stage data in SQS queues to inject metadata before accessing DynamoDB.
- D. Read records from DynamoDB and discard irrelevant data client-side.

Answer: B

NEW QUESTION 61

- (Exam Topic 1)

A Developer wants to use AWS X-Ray to trace a user request end-to-end throughout the software stack. The Developer made the necessary changes in the application tested it, and found that the application is able to send the traces to AWS X-Ray. However, when the application is deployed to an EC2 instance, the traces are not available.

Which of the following could create this situation? (Select two.)

- A. The traces are reaching X-Ray, but the Developer does not have access to view the records.
- B. The X-Ray daemon is not installed on the EC2 instance.
- C. The X-Ray endpoint specified in the application configuration is incorrect.
- D. The instance role does not have "xray:BatchGetTraces" and "xray:GetTraceGraph" permissions.
- E. The instance role does not have "xray:PutTraceSegments" and "xray:PutTelemetryRecords" permissions.

Answer: BE

NEW QUESTION 64

- (Exam Topic 1)

A Developer must build an application that uses Amazon DynamoDB. The requirements state that items being stored in the DynamoDB table will be 7KB in size and that reads must be strongly consistent. The maximum read rate is 3 items per second, and the maximum write rate is 10 items per second.

How should the Developer size the DynamoDB table to meet these requirements?

- A. Read: 3 read capacity units Write: 70 write capacity units
- B. Read: 6 read capacity units Write: 70 write capacity units
- C. Read: 6 read capacity units Write: 10 write capacity units
- D. Read: 3 read capacity units Write: 10 write capacity units

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Limits.html>

NEW QUESTION 65

- (Exam Topic 1)

An application will ingest data at a very high throughput from many sources and must store the data in an Amazon S3 bucket. Which service would BEST accomplish this task?

- A. Amazon Kinesis Firehose
- B. Amazon S3 Acceleration Transfer
- C. Amazon SQS
- D. Amazon SNS

Answer: A

NEW QUESTION 66

- (Exam Topic 1)

A Developer has created a large Lambda function, and deployment is failing with the following error: ClientError: An error occurred (InvalidParameterValueException) when calling the CreateFunction operation: Unzipped size must be smaller than XXXXXXXXX bytes', where XXXXXXXXX is the current Lambda limit

What can the Developer do to fix this problem?

- A. Submit a limit increase request to AWS Support to increase the function to the size needed.
- B. Use a compression algorithm that is more efficient than ZIP.
- C. Break the function into multiple smaller Lambda functions.
- D. ZIP the ZIP file twice to compress it further.

Answer: C

NEW QUESTION 71

- (Exam Topic 1)

A Developer is writing transactions into a DynamoDB table called "SystemUpdates" that has 5 write capacity units. Which option has the highest read throughput?

- A. Eventually consistent reads of 5 read capacity units reading items that are 4 KB in size
- B. Strongly consistent reads of 5 read capacity units reading items that are 4 KB in size
- C. Eventually consistent reads of 15 read capacity units reading items that are 1 KB in size
- D. Strongly consistent reads of 15 read capacity units reading items that are 1 KB in size

Answer: B

NEW QUESTION 72

- (Exam Topic 1)

A Developer is using AWS CLI, but when running list commands on a large number of resources, it is timing out. What can be done to avoid this time-out?

- A. Use pagination
- B. Use shorthand syntax
- C. Use parameter values
- D. Use quoting strings

Answer: A

Explanation:

Reference: <https://docs.aws.amazon.com/cli/latest/userguide/cli-usage-pagination.html>

NEW QUESTION 73

- (Exam Topic 1)

A customer wants to deploy its source code on an AWS Elastic Beanstalk environment. The customer needs to perform deployment with minimal outage and should only use existing instances to retain application access log.

What deployment policy would satisfy these requirements?

- A. Rolling
- B. All at once
- C. Rolling with an additional batch
- D. Immutable

Answer: A

NEW QUESTION 76

- (Exam Topic 1)

A company maintains a REST service using Amazon API Gateway and the API Gateway native API key validation. The company recently launched a new registration page, which allows users to sign up for the service. The registration page creates a new API key using CreateApiKey and sends the new key to the user. When the user attempts to call the API using this key, the user receives a 403 Forbidden error. Existing users are unaffected and can still call the API. What code updates will grant these new users access to the API?

- A. The createDeployment method must be called so the API can be redeployed to include the newly created API key.
- B. The updateAuthorizer method must be called to update the API's authorizer to include the newly created API key.
- C. The importApiKeys method must be called to import all newly created API keys into the current stage of the API.
- D. The createUsagePlanKey method must be called to associate the newly created API key with the correct usage plan.

Answer: D

Explanation:

<https://stackoverflow.com/questions/39061041/using-an-api-key-in-amazon-api-gateway>

NEW QUESTION 80

- (Exam Topic 1)

When a Simple Queue Service message triggers a task that takes 5 minutes to complete, which process below will result in successful processing of the message and remove it from the queue while minimizing the chances of duplicate processing?

- A. Retrieve the message with an increased visibility timeout, process the message, delete the message from the queue
- B. Retrieve the message with an increased visibility timeout, delete the message from the queue, process the message
- C. Retrieve the message with increased DelaySeconds, process the message, delete the message from the queue
- D. Retrieve the message with increased DelaySeconds, delete the message from the queue, process the message

Answer: A

Explanation:

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

NEW QUESTION 85

- (Exam Topic 1)

What AWS products and features can be deployed by Elastic Beanstalk? Choose 3 answers

- A. Auto scaling groups
- B. Route 53 hosted zones
- C. Elastic Load Balancers
- D. RDS Instances
- E. Elastic IP addresses
- F. SQS Queues

Answer: ACD

Explanation:

<https://aws.amazon.com/elasticbeanstalk/faqs/>

Q: What are the Cloud resources powering my AWS Elastic Beanstalk application? AWS Elastic Beanstalk uses proven AWS features and services, such as Amazon EC2, Amazon RDS, Elastic Load Balancing, Auto Scaling, Amazon S3, and Amazon SNS, to create an environment that runs your application. The current version of AWS Elastic Beanstalk uses the Amazon Linux AMI or the Windows Server 2012 R2 AMI.

NEW QUESTION 87

- (Exam Topic 1)

The release process workflow of an application requires a manual approval before the code is deployed into the production environment. What is the BEST way to achieve this using AWS CodePipeline?

- A. Use multiple pipelines to allow approval
- B. Use an approval action in a stage
- C. Disable the stage transition to allow manual approval
- D. Disable a stage just prior the deployment stage

Answer: B

NEW QUESTION 88

- (Exam Topic 1)

What type of block cipher does Amazon S3 offer for server side encryption?

- A. Triple DES
- B. Advanced Encryption Standard
- C. Blowfish
- D. RC5

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/serv-side-encryption.html>

NEW QUESTION 91

- (Exam Topic 1)

Which of the following are valid SNS delivery transports? Choose 2 answers

- A. HTTP
- B. UDP
- C. SMS
- D. DynamoDB
- E. Named Pipes

Answer: AC

Explanation:

<https://aws.amazon.com/sns/faqs/>

Q: What are the different delivery formats/transport for receiving notifications?

In order for customers to have broad flexibility of delivery mechanisms, Amazon SNS supports notifications over multiple transport protocols. Customers can select one of the following transports as part of the subscription requests:

“HTTP”, “HTTPS” – Subscribers specify a URL as part of the subscription registration; notifications will be delivered through an HTTP POST to the specified URL.
“Email”, “Email-JSON” – Messages are sent to registered addresses as email. Email-JSON sends notifications as a JSON object, while Email sends text-based email.

“SQS” – Users can specify an SQS standard queue as the endpoint; Amazon SNS will enqueue a notification message to the specified queue (which subscribers can then process using SQS APIs such as ReceiveMessage, DeleteMessage, etc.). Note that FIFO queues are not currently supported.

“SMS” – Messages are sent to registered phone numbers as SMS text messages.

NEW QUESTION 96

- (Exam Topic 1)

If an application is storing hourly log files from thousands of instances from a high traffic web site, which naming scheme would give optimal performance on S3?

- A. Sequential
- B. instanceID_log-HH-DD-MM-YYYY
- C. instanceID_log-YYYY-MM-DD-HH
- D. HH-DD-MM-YYYY-log_instanceID
- E. YYYY-MM-DD-HH-log_instanceID

Answer: B

Explanation:

Reference:

https://acloud.guru/forums/aws-certified-developer-associate/discussion/-KU2dEtJb-LI5ISbH_S4/if-an-applicati

NEW QUESTION 98

- (Exam Topic 1)

A current architecture uses many Lambda functions invoking one another as a large state machine. The coordination of this state machine is legacy custom code that breaks easily.

Which AWS Service can help refactor and manage the state machine?

- A. AWS Data Pipeline
- B. AWS SNS with AWS SQS
- C. Amazon Elastic MapReduce
- D. AWS Step Functions

Answer: D

Explanation:

<https://aws.amazon.com/step-functions/>

NEW QUESTION 102

- (Exam Topic 1)

The Developer for a retail company must integrate a fraud detection solution into the order processing solution. The fraud detection solution takes between ten and thirty minutes to verify an order. At peak, the web site can receive one hundred orders per minute.

What is the most scalable method to add the fraud detection solution to the order processing pipeline?

- A. Add all new orders to an Amazon SQS queue
- B. Configure a fleet of 10 EC2 instances spanning multiple AZs with the fraud detection solution installed on them to pull orders from this queue
- C. Update the order with a pass or fails status.
- D. Add all new orders to an SQS queue
- E. Configure an Auto Scaling group that uses the queue depth metric as its unit of scale to launch a dynamically-sized fleet of EC2 instances spanning multiple AZs with the fraud detection solution installed on them to pull orders from this queue
- F. Update the order with a pass or fails status.
- G. Add all new orders to an Amazon Kinesis Stream
- H. Subscribe a Lambda function to automatically read batches of records from the Kinesis Stream
- I. The Lambda function includes the fraud detection software and will update the order with a pass or fail status.
- J. Write all new orders to Amazon DynamoDB
- K. Configure DynamoDB Streams to include all new orders. Subscribe a Lambda function to automatically read batches of records from the Kinesis Stream
- L. The Lambda function includes the fraud detection software and will update the order with a pass or fail status.

Answer: B

NEW QUESTION 103

- (Exam Topic 1)

Games-R-Us is launching a new game app for mobile devices. Users will log into the game using their existing Facebook account and the game will record player data and scoring information directly to a DynamoDB table.

What is the most secure approach for signing requests to the DynamoDB API?

- A. Create an IAM user with access credentials that are distributed with the mobile app to sign the requests
- B. Distribute the AWS root account access credentials with the mobile app to sign the requests
- C. Request temporary security credentials using web identity federation to sign the requests
- D. Establish cross account access between the mobile app and the DynamoDB table to sign the requests

Answer: C

NEW QUESTION 104

- (Exam Topic 1)

A company needs to encrypt data at rest, but it wants to leverage an AWS managed service using its own master key. Which of the following AWS service can be used to meet these requirements?

- A. SSE with Amazon S3
- B. SSE with AWS KMS
- C. Client-side encryption
- D. AWS IAM roles and policies

Answer: B

NEW QUESTION 108

- (Exam Topic 1)

How should custom libraries be utilized in AWS Lambda?

- A. Host the library on Amazon S3 and reference to it from the Lambda function.
- B. Install the library locally and upload a ZIP file of the Lambda function.
- C. Import the necessary Lambda blueprint when creating the function.
- D. Modify the function runtime to include the necessary library.

Answer: D

Explanation:

Reference: https://docs.aws.amazon.com/lambda/latest/dg/env_variables.html

NEW QUESTION 110

- (Exam Topic 1)

A Developer is writing an imaging micro service on AWS Lambda. The service is dependent on several libraries that are not available in the Lambda runtime environment.

Which strategy should the Developer follow to create the Lambda deployment package?

- A. Create a ZIP file with the source code and all dependent libraries.
- B. Create a ZIP file with the source code and a script that installs the dependent libraries at runtime.
- C. Create a ZIP file with the source code
- D. Stage the dependent libraries on an Amazon S3 bucket indicated by the Lambda environment variable LD_LIBRARY_PATH
- E. Create a ZIP file with the source code and a buildspec.yaml file that installs the dependent libraries on AWS Lambda.

Answer: B

NEW QUESTION 112

- (Exam Topic 1)

A Developer has written a serverless application using multiple AWS services. The business logic is written as a Lambda function which has dependencies on third-party libraries. The Lambda function endpoints will be exposed using Amazon API Gateway. The Lambda function will write the information to Amazon DynamoDB. The Developer is ready to deploy the application but must have the ability to rollback. How can this deployment be automated, based on these requirements?

- A. Deploy using Amazon Lambda API operations to create the Lambda function by providing a deployment package.
- B. Use an AWS CloudFormation template and use CloudFormation syntax to define the Lambda function resource in the template.
- C. Use syntax conforming to the Serverless Application Model in the AWS CloudFormation template to define the Lambda function resource.
- D. Create a bash script which uses AWS CLI to package and deploy the application.

Answer: C

Explanation:

Refer AWS documentation - SAM Gradual Code Deployment

If you use AWS SAM to create your serverless application, it comes built-in with AWS CodeDeploy to help ensure safe Lambda deployments. With just a few lines of configuration, AWS SAM does the following for you:

- Deploys new versions of your Lambda function, and automatically creates aliases that point to the new version.
- Gradually shifts customer traffic to the new version until you're satisfied that it's working as expected, or you roll back the update.
- Defines pre-traffic and post-traffic test functions to verify that the newly deployed code is configured correctly and your application operates as expected.
- Rolls back the deployment if CloudWatch alarms are triggered.

NEW QUESTION 113

- (Exam Topic 1)

A Developer has been asked to build a real-time dashboard web application to visualize the key prefixes and storage size of objects in Amazon S3 buckets. Amazon DynamoDB will be used to store the Amazon S3 metadata.

What is the optimal and MOST cost-effective design to ensure that the real-time dashboard is kept up to date with the state of the objects in the Amazon S3 buckets?

- A. Use an Amazon CloudWatch event backed by an AWS Lambda function
- B. Issue an Amazon S3 API call to get a list of all Amazon S3 objects and persist the metadata within DynamoDB
- C. Have the web application poll the DynamoDB table to reflect this change.
- D. Use Amazon S3 Event Notification backed by a Lambda function to persist the metadata into DynamoDB
- E. Have the web application poll the DynamoDB table to reflect this change.
- F. Run a cron job within an Amazon EC2 instance to list all objects within Amazon S3 and persist the metadata into DynamoDB
- G. Have the web application poll the DynamoDB table to reflect this change.
- H. Create a new Amazon EMR cluster to get all the metadata about Amazon S3 objects; persist the metadata into DynamoDB

I. Have the web application poll the DynamoDB table to reflect this change.

Answer: A

NEW QUESTION 115

- (Exam Topic 1)

Which of the following services are included at no additional cost with the use of the AWS platform? Choose 2 answers

- A. Simple Storage Service
- B. Elastic Compute Cloud
- C. Auto Scaling
- D. Elastic Load Balancing
- E. CloudFormation
- F. Simple Workflow Service

Answer: CE

NEW QUESTION 120

- (Exam Topic 1)

A web application is using Amazon Kinesis 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 Streams?

- A. Enable SSL connections to Kinesis
- B. Use Amazon Kinesis Consumer Library
- C. Encrypt the data once it is at rest with a Lambda function
- D. Enable server-side encryption in Kinesis Streams

Answer: D

Explanation:

<https://docs.aws.amazon.com/streams/latest/dev/what-is-sse.html>

<https://aws.amazon.com/about-aws/whats-new/2017/07/amazon-kinesis-streams-introduces-server-side-encrypti>

NEW QUESTION 122

- (Exam Topic 1)

A Development team currently supports an application that uses an in-memory store to save accumulated game results. Individual results are stored in a database. As part of migrating to AWS, the team needs to use automatic scaling. The team knows this will yield inconsistent results. Where should the team store these accumulated game results to BEST allow for consistent results without impacting performance?

- A. Amazon S3
- B. Amazon RDS
- C. Amazon ElastiCache
- D. Amazon Kinesis

Answer: C

NEW QUESTION 123

- (Exam Topic 1)

You are providing AWS consulting services for a company developing a new mobile application that will be leveraging Amazon SNS Mobile Push for push notifications. In order to send direct notification messages to individual devices each device registration identifier or token needs to be registered with SNS; however the developers are not sure of the best way to do this. You advise them to:

- A. Bulk upload the device tokens contained in a CSV file via the AWS Management Console.
- B. Let the push notification service (e.g. Amazon Device Messaging) handle the registration.
- C. Implement a token vending service to handle the registration.
- D. Call the CreatePlatformEndPoint API function to register multiple device tokens.

Answer: D

Explanation:

<https://docs.aws.amazon.com/sns/latest/dg/mobile-push-send-devicetoken.html>

NEW QUESTION 128

- (Exam Topic 1)

An application is designed to use Amazon SQS to manage messages from many independent senders. Each sender's messages must be processed in the order they are received.

Which SQS feature should be implemented by the Developer?

- A. Configure each sender with a unique MessageGroupId
- B. Enable MessageDeduplicationIds on the SQS queue
- C. Configure each message with unique MessageGroupIds.
- D. Enable ContentBasedDeduplication on the SQS queue

Answer: A

NEW QUESTION 133

- (Exam Topic 1)

An application on AWS is using third-party APIs. The Developer needs to monitor API errors in the code, and wants to receive notifications if failures go above a set threshold value.

How can the Developer achieve these requirements?

- A. Publish a custom metric on Amazon CloudWatch and use Amazon SES for notification.
- B. Use an Amazon CloudWatch API-error metric and use Amazon SNS for notification.
- C. Use an Amazon CloudWatch API-error metric and use Amazon SES for notification.
- D. Publish a custom metric on Amazon CloudWatch and use Amazon SNS for notification.

Answer: D

NEW QUESTION 136

- (Exam Topic 1)

A Developer is asked to implement a caching layer in front of Amazon RDS. Cached content is expensive to regenerate in case of service failure. Which implementation below would work while maintaining maximum uptime?

- A. Implement Amazon ElastiCache Redis in Cluster Mode
- B. Install Redis on an Amazon EC2 instance.
- C. Implement Amazon ElastiCache Memcached.
- D. Migrate the database to Amazon Redshift.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/SelectEngine.html>

NEW QUESTION 140

- (Exam Topic 1)

A company has multiple Developers located across the globe who are updating code incrementally for a development project. When Developers upload code concurrently, internet connectivity is slow, and it is taking a long time to upload code for deployment in AWS Elastic Beanstalk.

Which step will result in minimized upload and deployment time with the LEAST amount of administrative effort?

- A. Allow the Developers to upload the code to an Amazon S3 bucket, and deploy it directly to Elastic Beanstalk.
- B. Allow the Developers to upload the code to a central FTP server to deploy the application to Elastic Beanstalk.
- C. Create an AWS CodeCommit repository, allow the Developers to commit code to it, and then directly deploy the code to Elastic Beanstalk.
- D. Create a code repository on an Amazon EC2 instance so that all Developers can update the code, and deploy the application from the instance to Elastic Beanstalk.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/deploy-codecommit-elastic-beanstalk/>

NEW QUESTION 143

- (Exam Topic 1)

A Lambda function is packaged for deployment to multiple environments, including development, test, production, etc. Each environment has unique set of resources such as databases, etc.

How can the Lambda function use the resources for the current environment?

- A. Apply tags to the Lambda functions.
- B. Hardcore resources in the source code.
- C. Use environment variables for the Lambda functions.
- D. Use separate function for development and production.

Answer: C

NEW QUESTION 146

- (Exam Topic 1)

A Developer created configuration specifications for an AWS Elastic Beanstalk application in a file named healthcheckurl.yaml in the .ebextensions/ directory of their application source bundle. The file contains the following:

```
option_settings:
  - namespace: aws:elasticbeanstalk:application
    option_name: Application Healthcheck URL
    value: /health_check
```

After the application launches, the health check is not being run on the correct path, even though it is valid. What can be done to correct this configuration file?

- A. Convert the file to JSON format.
- B. Rename the file to a .config extension.
- C. Change the configuration section from options_settings to resources.

D. Change the namespace of the option settings to a custom namespace.

Answer: B

Explanation:

Reference: <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/ebextensions.html>

You can add AWS Elastic Beanstalk configuration files (.ebextensions) to your web application's source code to configure your environment and customize the AWS resources that it contains. Configuration files are YAML- or JSON-formatted documents with a .config file extension that you place in a folder named .ebextensions and deploy in your application source bundle. <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/ebextensions.html>

NEW QUESTION 148

- (Exam Topic 1)

An AWS Lambda function generates a 3MB JSON file and then uploads it to an Amazon S3 bucket daily. The file contains sensitive information, so the Developer must ensure that it is encrypted before uploading to the bucket.

Which of the following modifications should the Developer make to ensure that the data is encrypted before uploading it to the bucket?

- A. Use the default AWS KMS customer master key for S3 in the Lambda function code.
- B. Use the S3 managed key and call the GenerateDataKey API to encrypt the file.
- C. Use the GenerateDataKey API, then use that data key to encrypt the file in the Lambda function code.
- D. Use a custom KMS customer master key created for S3 in the Lambda function code.

Answer: C

NEW QUESTION 153

- (Exam Topic 1)

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 IAM 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: D

Explanation:

<https://docs.aws.amazon.com/cognito/latest/developerguide/iam-roles.html>

NEW QUESTION 155

- (Exam Topic 1)

A Developer needs to use AWS X-Ray to monitor an application that is deployed on EC2 instances. What steps have to be executed to perform the monitoring?

- A. Deploy the X-Ray SDK with the application and use X-Ray annotation.
- B. Install the X-Ray daemon and instrument the application code.
- C. Install the X-Ray daemon and configure it to forward data to Amazon CloudWatch Events.
- D. Deploy the X-Ray SDK with the application and instrument the application code.

Answer: C

NEW QUESTION 159

- (Exam Topic 1)

A nightly batch job loads 1 million new records into a DynamoDB table. The records are only needed for one hour, and the table needs to be empty by the next night's batch job.

Which is the MOST efficient and cost-effective method to provide an empty table?

- A. Use DeleteItem using a ConditionExpression.
- B. Use BatchWriteItem to empty all of the rows.
- C. With a recursive function that scans and calls out DeleteItem.
- D. Create and then delete the table after the task has completed.

Answer: D

Explanation:

"Deleting an entire table is significantly more efficient than removing items one-by-one, which essentially doubles the write throughput as you do as many delete operations as put operations"

NEW QUESTION 163

- (Exam Topic 1)

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 Glue to perform frequent ETL updates 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 Amazon Kinesis Data Firehose to deliver all changes from the Accounts database to the Payments database.
- D. Use Amazon DynamoDB Streams to deliver all changes from the Accounts database to the Payments database.

Answer:

D

Explanation:

Reference:

<https://aws.amazon.com/blogs/database/how-to-perform-ordered-data-replication-between-applications-by-using>**NEW QUESTION 167**

- (Exam Topic 1)

An application takes 40 seconds to process instructions received in an Amazon SQS message.

Assuming the SQS queue is configured with the default VisibilityTimeout value, what is the BEST way, upon receiving a message, to ensure that no other instances can retrieve a message that has already been processed or is currently being processed?

- A. Use the ChangeMessageVisibility API to increase the VisibilityTimeout, then use the DeleteMessage API to delete the message.
- B. Use the DeleteMessage API call to delete the message from the queue, then call DeleteQueue API to remove the queue.
- C. Use the ChangeMessageVisibility API to decrease the timeout value, then use the DeleteMessage API to delete the message.
- D. Use the DeleteMessageVisibility API to cancel the VisibilityTimeout, then use the DeleteMessage API to delete the message.

Answer: A**Explanation:**<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html> In SQS, messages remain there. It is the consumer's responsibility to delete it, once consumed and processed.**NEW QUESTION 172**

- (Exam Topic 1)

A company is developing an application that will run on several Amazon EC2 instances in an Auto Scaling

group and can access a database running on Amazon EC2. The application needs to store secrets required to connect to the database. The application must allow for periodic secret rotation, and there should be no changes to the application when a secret changes.

What is the SAFEST way to meet these requirements?

- A. Associate an IAM role to the EC2 instance where the application is running with permission to access the database.
- B. Use AWS Systems Manager Parameter Store with the SecureString data type to store secrets.
- C. Configure the application to store secrets in Amazon S3 object metadata.
- D. Hard code the database secrets in the application code itself.

Answer: B**NEW QUESTION 177**

- (Exam Topic 1)

A Developer created a dashboard for an application using Amazon API Gateway, Amazon S3, AWS Lambda, and Amazon RDS. The Developer needs an authentication mechanism allowing a user to sign in and view the dashboard. It must be accessible from mobile applications, desktops, and tablets, and must remember user preferences across platforms.

Which AWS service should the Developer use to support this authentication scenario?

- A. AWS KMS
- B. Amazon Cognito
- C. AWS Directory Service
- D. Amazon IAM

Answer: B**Explanation:**

Cognito user pool provides sign up and sign in functionality along with identity pool which provides temp credentials for using aws services.

NEW QUESTION 181

- (Exam Topic 1)

A company wants to migrate its web application to AWS and leverage Auto Scaling to handle peak workloads. The Solutions Architect determined that the best metric for an Auto Scaling event is the number of concurrent users.

Based on this information, what should the Developer use to autoscale based on concurrent users?

- A. An Amazon SNS topic to be triggered when a concurrent user threshold is met
- B. An Amazon CloudWatch Network metric
- C. Amazon CloudFront to leverage AWS Edge Locations
- D. A Custom Amazon CloudWatch metric for concurrent users.

Answer: D**NEW QUESTION 184**

- (Exam Topic 1)

Which DynamoDB limits can be raised by contacting AWS support? Choose 2 answers

- A. The number of hash keys per account
- B. The maximum storage used per account
- C. The number of tables per account
- D. The number of local secondary indexes per account
- E. The number of provisioned throughput units per account

Answer: CE

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Limits.html>

NEW QUESTION 185

- (Exam Topic 1)

A Developer executed a AWS CLI command and received the error shown below:

```
A client error (UnauthorizedOperation) occurred when calling the RunInstances operation:
You are not authorized to perform this operation. Encoded authorization failure message:
oGsbAaIV7wlfj8zUqebHUANHzFbmkzILlxyj_y9xwhIhk99U_cUq1FIeZnskWDjQ1wSHStVfdCEyZILGocccGpCiC
IhORceWF9rRwFTnEcRj3N9iTrPAE1WHveC5Z54ALPaWlEjHlLg8CaB8d8lCKmxQuylCm0r1Bf2fHJRujAYopMvmga
8olFmKAl9yn_25rI120Q9p5ZIMX28zYM4dTulcJQUQjosgrEejfiIMYDda8l7Ooko9H6VmGJX62KfkRa5l7yE6hhh
2bIwA6tpyCJy2LWFRTe4bafqAyoqkarhPA4mGiZyWn4gSqbO8o-
uqSIvWYPweaKGkampa0arcFR4gBD7Ph097WYBkzX9hVjGppLMY4jpxRv
```

What action should the Developer perform to make this error human-readable?

- A. Make a call to AWS KMS to decode the message.
- B. Use the AWS STS decode-authorization-message API to decode the message.
- C. Use an open source decoding library to decode the message.
- D. Use the AWS IAM decode-authorization-message API to decode this message.

Answer: B

Explanation:

<https://docs.aws.amazon.com/cli/latest/reference/sts/decode-authorization-message.html>

The message is encoded because the details of the authorization status can constitute privileged information that the user who requested the operation should not see. To decode an authorization status message, a user must be granted permissions via an IAM policy to request the `DecodeAuthorizationMessage` (`sts:DecodeAuthorizationMessage`) action.

NEW QUESTION 188

- (Exam Topic 1)

A company has written a Java AWS Lambda function to be triggered whenever a user uploads an image to an Amazon S3 bucket. The function converts the original image to several different formats and then copies the resulting images to another Amazon S3 bucket.

The Developers find that no images are being copied to the second Amazon S3 bucket. They have tested the code on an Amazon EC2 instance with 1GB of RAM, and it takes an average of 500 seconds to complete.

What is the MOST likely cause of the problem?

- A. The Lambda function has insufficient memory and needs to be increased to 1 GB to match the Amazon EC2 instance
- B. Files need to be copied to the same Amazon S3 bucket for processing, so the second bucket needs to be deleted.
- C. Lambda functions have a maximum execution limit of 300 seconds, therefore the function is not completing.
- D. There is a problem with the Java runtime for Lambda, and the function needs to be converted to node.js.

Answer: C

NEW QUESTION 193

- (Exam Topic 1)

An application overwrites an object in Amazon S3, and then immediately reads the same object. Why would the application sometimes retrieve the old version of the object?

- A. S3 overwrite PUTS are eventually consistent, so the application may read the old object.
- B. The application needs to add extra metadata to label the latest version when uploading to Amazon S3.
- C. All S3 PUTS are eventually consistent, so the application may read the old object.
- D. The application needs to explicitly specify latest version when retrieving the object.

Answer: A

NEW QUESTION 195

- (Exam Topic 1)

A company is providing services to many downstream consumers. Each consumer may connect to one or more services. This has resulted in a complex architecture that is difficult to manage and does not scale well. The company needs a single interface to manage these services to consumers.

Which AWS service should be used to refactor this architecture?

- A. AWS Lambda
- B. AWS X-Ray
- C. Amazon SQS
- D. Amazon API Gateway

Answer: D

NEW QUESTION 200

- (Exam Topic 1)

A set of APIs are exposed to customers using the Amazon API Gateway. These APIs have caching enabled on the API Gateway. Customers have asked for an option to invalidate this cache for each of the APIs.

What action can be taken to allow API customers to invalidate the API Cache?

- A. Ask customers to use AWS credentials to call the `InvalidateCache` API.
- B. Ask customers to invoke an AWS API endpoint which invalidates the cache.

- C. Ask customers to pass an HTTP header called Cache-Control:max-age=0.
- D. Ask customers to add a query string parameter called "INVALIDATE_CACHE" when making an API call.

Answer: C

Explanation:

<https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Cache-Control> <https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-caching.html>

Invalidate an API Gateway Cache Entry A client of your API can invalidate an existing cache entry and reload it from the integration endpoint for individual requests. The client must send a request that contains the Cache-Control: max-age=0 header. The client receives the response directly from the integration endpoint instead of the cache, provided that the client is authorized to do so. This replaces the existing cache entry with the new response, which is fetched from the integration endpoint.

NEW QUESTION 205

- (Exam Topic 1)

A company is migrating its on-premises database to Amazon RDS for MySQL. The company has read-heavy workloads, and wants to make sure it re-factors its code to achieve optimum read performance for its queries.

How can this objective be met?

- A. Add database retries to effectively use RDS with vertical scaling
- B. Use RDS with multi-AZ deployment
- C. Add a connection string to use an RDS read replica for read queries
- D. Add a connection string to use a read replica on an EC2 instance.

Answer: C

NEW QUESTION 206

- (Exam Topic 1)

A development team consists of 10 team members. Similar to a home directory for each team member the manager wants to grant access to user-specific folders in an Amazon S3 bucket. For the team member with the username "TeamMemberX", the snippet of the IAM policy looks like this:

```
{ "Sid": "AllowS3ActionToFolders", "Effect": "Allow", "Action":  
["s3:*"], "Resource":  
["arn:aws:s3:::companyname/home/TeamMemberX/*"] }
```

Instead of creating distinct policies for each team member, what approach can be used to make this policy snippet generic for all team members?

- A. Use IAM policy condition
- B. Use IAM policy principal
- C. Use IAM policy variables
- D. Use IAM policy resource

Answer: C

Explanation:

> https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_variables.html > Use AWS Identity and Access Management (IAM) policy variables as placeholders when you don't know the exact value of a resource or condition key when you write the policy.

NEW QUESTION 207

- (Exam Topic 1)

A serverless application uses an API Gateway and AWS Lambda.

Where should the Lambda function store its session information across function calls?

- A. In an Amazon DynamoDB table
- B. In an Amazon SQS queue
- C. In the local filesystem
- D. In an SQLite session table using `-DSQLITE_ENABLE_SESSION`

Answer: A

NEW QUESTION 209

- (Exam Topic 1)

A Developer is creating a web application that requires authentication, but also needs to support guest access to provide users limited access without having to authenticate. What service can provide support for the application to allow guest access?

- A. IAM temporary credentials using AWS STS.
- B. Amazon Directory Service
- C. Amazon Cognito with unauthenticated access enabled
- D. IAM with SAML integration

Answer: C

Explanation:

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-getting-started-hello> <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/sam-cli-command-reference-sa>

NEW QUESTION 211

- (Exam Topic 1)

A company is migrating a single-server, on-premises web application to AWS. The company intends to use multiple servers behind an Elastic Load Balancer (ELB) to balance the load, and will also store session data in memory on the web server. The company does not want to lose that session data if a server fails or goes offline, and it wants to minimize user's downtime.

Where should the company move session data to MOST effectively reduce downtime and make users' session data more fault tolerant?

- A. An Amazon ElastiCache for Redis cluster
- B. A second Amazon EBS volume
- C. The web server's primary disk
- D. An Amazon EC2 instance dedicated to session data

Answer: A

NEW QUESTION 213

- (Exam Topic 1)

A Developer has implemented a Lambda function that needs to add new customers to an RDS database that is expected to run hundreds of times per hour. The Lambda function is configured to use 512MB of RAM and is based on the following pseudo code:

```
def lambda_handler(event, context):
```

```
    db = database.connect()
```

```
    db.statement('INSERT INTO Customers (CustomerName) VALUES  
                (context.name)')
```

```
    db.close()
```

After testing the Lambda function, the Developer notices that the Lambda execution time is much longer than expected. What should the Developer do to improve performance?

- A. Increase the amount of RAM allocated to the Lambda function, which will increase the number of threads the Lambda can use.
- B. Increase the size of the RDS database to allow for an increased number of database connections each hour.
- C. Move the database connection and close statement out of the handle
- D. Place the connection in the global space.
- E. Replace RDS with Amazon DynamoDB to implement control over the number of writes per second.

Answer: C

Explanation:

Refer AWS documentation - Lambda Best Practices

Take advantage of Execution Context reuse to improve the performance of your function. Make sure any externalized configuration or dependencies that your code retrieves are stored and referenced locally after initial execution. Limit the re-initialization of variables/objects on every invocation. Instead use static initialization/constructor, global/static variables and singletons. Keep alive and reuse connections (HTTP, database, etc.) that were established during a previous invocation.

NEW QUESTION 214

- (Exam Topic 1)

A supplier is writing a new RESTful API for customers to query the status of orders. The customers requested the following API endpoint.

<http://www.supplierdomain.com/status/customerID>

Which of the following application designs meet the requirements? (Select two.)

- A. Amazon SQS; Amazon SNS
- B. Elastic Load Balancing; Amazon EC2
- C. Amazon ElastiCache; Amazon Elasticservice
- D. Amazon API Gateway; AWS Lambda
- E. Amazon S3; Amazon CloudFront

Answer: DE

NEW QUESTION 217

- (Exam Topic 1)

A large e-commerce site is being designed to deliver static objects from Amazon S3. The Amazon S3 bucket will server more than 300 GET requests per second. What should be done to optimize performance? (Select TWO.)

- A. Integrate Amazon CloudFront with Amazon S3.
- B. Enable Amazon S3 cross-region replication.
- C. Delete expired Amazon S3 server log files.
- D. Configure Amazon S3 lifecycle rules.
- E. Randomize Amazon S3 key name prefixes.

Answer: AE

Explanation:

CloudWatch definitely. Random key prefixes is still a valid method of improving performance by using parallel reads. It doesn't mention prefix hashing. For instance prefixes 1/,2/,3/,4/,5/ could provide 5 x parallel streams for S3 as opposed to all objects being in a single folder/prefix e.g. dev/
<https://docs.aws.amazon.com/AmazonS3/latest/dev/optimizing-performance.html>

"There are no limits to the number of prefixes in a bucket. You can increase your read or write performance by parallelizing reads. For example, if you create 10

prefixes in an Amazon S3 bucket to parallelize reads, you could scale your read performance to 55,000 read requests per second." The assumption that prefixes don't matter is incorrect, as described by "Amazon S3 performance guidelines recommended randomizing prefix naming with ****hashed characters**** to optimize performance for frequent data retrievals. You no longer have to randomize prefix naming for performance, and can use sequential date-based naming for your prefixes"

NEW QUESTION 219

- (Exam Topic 1)

A company is building an application to track athlete performance using an Amazon DynamoDB table. Each item in the table is identified by a partition key (user_id) and a sort key (sport_name). The table design is shown below: (Note: Not all table attributes are shown)

A Developer is asked to write a leaderboard application to display the top performers (user_id) based on the score for each sport_name.

What process will allow the Developer to extract results MOST efficiently from the DynamoDB table?

- A. Use a DynamoDB query operation with the key attributes of user_id and sport_name and order the results based on the score attribute.
- B. Create a global secondary index with a partition key of sport_name and a sort key of score, and get the results
- C. Use a DynamoDB scan operation to retrieve scores and user_id based on sport_name, and order the results based on the score attribute.
- D. Create a local secondary index with a primary key of sport_name and a sort key of score and get the results based on the score attribute.

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html>

https://docs.aws.amazon.com/zh_cn/amazondynamodb/latest/developerguide/GSI.html

NEW QUESTION 224

- (Exam Topic 1)

An application that runs on an Amazon EC2 instance needs to access and make API calls to multiple AWS services.

What is the MOST secure way to provide access to the AWS services with MINIMAL management overhead?

- A. Use AWS KMS to store and retrieve credentials.
- B. Use EC2 instance profiles.
- C. Use AWS root user to make requests to the application.
- D. Store and retrieve credentials from AWS CodeCommit.

Answer: B

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html

NEW QUESTION 228

- (Exam Topic 1)

A meteorological system monitors 600 temperature gauges, obtaining temperature samples every minute and saving each sample to a DynamoDB table. Each sample involves writing 1K of data and the writes are evenly distributed over time.

How much write throughput is required for the target table?

- A. 1 write capacity unit
- B. 10 write capacity units
- C. 60 write capacity units
- D. 600 write capacity units
- E. 3600 write capacity units

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h>

NEW QUESTION 230

- (Exam Topic 1)

A deployment package uses the AWS CLI to copy files into any S3 bucket in the account, using access keys stored in environment variables. The package is running on EC2 instances, and the instances have been modified to run with an assumed IAM role and a more restrictive policy that allows access to only one bucket.

After the change, the Developer logs into the host and still has the ability to write into all of the S3 buckets in that account.

What is the MOST likely cause of this situation?

- A. An IAM inline policy is being used on the IAM role
- B. An IAM managed policy is being used on the IAM role
- C. The AWS CLI is corrupt and needs to be reinstalled
- D. The AWS credential provider looks for instance profile credentials last

Answer: B

Explanation:

<https://docs.aws.amazon.com/sdk-for-java/v1/developer-guide/credentials.html>

NEW QUESTION 232

- (Exam Topic 1)

When a Developer tries to run an AWS CodeBuild project, it raises an error because the length of all environment variables exceeds the limit for the combined maximum of characters.

What is the recommended solution?

- A. Add the export LC_ALL="en_US.utf8" command to the pre_build section to ensure POSIX localization.
- B. Use Amazon Cognito to store key-value pairs for large numbers of environment variables.
- C. Update the settings for the build project to use an Amazon S3 bucket for large numbers of environment variables.
- D. Use AWS Systems Manager Parameter Store to store large numbers of environment variables.

Answer: D

NEW QUESTION 236

- (Exam Topic 1)

A Developer is creating a Lambda function and will be using external libraries that are not included in the standard Lambda libraries. What action would minimize the Lambda compute time consumed?

- A. Install the dependencies and external libraries at the beginning of the Lambda function.
- B. Create a Lambda deployment package that includes the external libraries.
- C. Copy the external libraries to Amazon S3, and reference the external libraries to the S3 location.
- D. Install the external libraries in Lambda to be available to all Lambda functions.

Answer: D

NEW QUESTION 240

- (Exam Topic 2)

A development team is designing a mobile app that requires multi-factor authentication Which steps should be taken to achieve this? (Select TWO)

- A. Use Amazon Cognito to create a user pool and create users in the user pool
- B. Send multi-factor authentication text codes to users with the Amazon SNS Publish API call in the app code
- C. Enable multi-factor authentication for the Amazon Cognito user pool
- D. Use AWS IAM to create IAM users
- E. Enable multi-factor authentication for the users created in AWS IAM.

Answer: AC

NEW QUESTION 245

- (Exam Topic 2)

A developer added a new feature to an application running on an Amazon EC2 instance that uses Amazon SQS After deployment, the developer noticed a significant increase in Amazon SQS costs. When monitoring the Amazon SQS metrics on Amazon CloudWatch. the developer found that on average one message per minute is posted on this queue.

What can be done to reduce Amazon SQS costs for this application?

- A. Increase the Amazon SQS queue polling timeout
- B. Scale down the Amazon SQS queue to the appropriate size for low traffic demand.
- C. Configure push delivery via Amazon SNS instead of polling the Amazon SQS queue
- D. Use an Amazon SQS first-in, first-out (FIFO) queue instead of a standard queue.

Answer: A

NEW QUESTION 247

- (Exam Topic 2)

A company has an application where reading objects from Amazon S3 is based on the type of user The user types are registered user and guest user The company has 25.000 users and is growing Information is pulled from an S3 bucket depending on the user type.

Which approaches are recommended to provide access to both user types? (Select TWO.)

- A. Provide a different access key and secret access key in the application code for registered users and guest users to provide read access to the objects
- B. Use S3 bucket policies to restrict read access to specific IAM users
- C. Use Amazon Cognito to provide access using authenticated and unauthenticated roles
- D. Create a new IAM user for each user and grant read access.
- E. Use the AWS IAM service and let the application assume the different roles using the AWS Security Token Service (AWS STS) AssumeRole action depending on the type of user and provide read access to Amazon S3 using the assumed role

Answer: BC

NEW QUESTION 251

- (Exam Topic 2)

A Developer is building a three-tier web application that should be able to handle a minimum of 5000 requests per minute. Requirements state that the web tier should be completely stateless while the application maintains session state for the users.

How can session data be externalized, keeping latency at the LOWEST possible value?

- A. Create an Amazon RDS instance, then implement session handling at the application level to leverage a database inside the RDS database instance for session data storage
- B. Implement a shared file system solution across the underlying Amazon EC2 instances, then implement session handling at the application level to leverage the shared file system for session data storage
- C. Create an Amazon ElastiCache Memcached cluster, then implement session handling at the application level to leverage the cluster for session data storage
- D. Create an Amazon DynamoDB table, then implement session handling at the application level to leverage the table for session data storage

Answer: A

Explanation:

<https://aws.amazon.com/caching/session-management/>

NEW QUESTION 255

- (Exam Topic 2)

A Developer is going to deploy an AWS Lambda function that requires significant CPU utilization. Which approach will MINIMIZE the average runtime of the function?

- A. Deploy the function into multiple AWS Regions
- B. Deploy the function into multiple Availability Zones
- C. Deploy the function using Lambda layers
- D. Deploy the function with its memory allocation set to the maximum amount

Answer: D

NEW QUESTION 259

- (Exam Topic 2)

An application needs to encrypt data that is written to Amazon S3 where the keys are managed in an on-premises data center and the encryption is handled by S3. Which type of encryption should be used?

- A. Use server-side encryption with Amazon S3-managed keys
- B. Use server-side encryption with AWS KMS-managed keys.
- C. Use client-side encryption with customer master keys
- D. Use server-side encryption with customer-provided keys

Answer: A

NEW QUESTION 263

- (Exam Topic 2)

A Developer wants to encrypt new objects that are being uploaded to an Amazon S3 bucket by an application. There must be an audit trail of who has used the key during this process. There should be no change to the performance of the application. Which type of encryption meets these requirements?

- A. Server-side encryption using S3-managed keys
- B. Server-side encryption with AWS KMS-managed keys
- C. Client-side encryption with a client-side symmetric master key
- D. Client-side encryption with AWS KMS-managed keys

Answer: B

NEW QUESTION 266

- (Exam Topic 2)

A development team is using AWS Elastic Beanstalk to deploy a two-tier application that consists of a load-balanced web tier and an Amazon RDS database tier in production. The team would like to separate the RDS instance from the Elastic Beanstalk. How can this be accomplished?

- A. Use the Elastic Beanstalk CLI to disassociate the database.
- B. Use the AWS CLI to disassociate the database.
- C. Change the deployment policy to disassociate the database.
- D. Recreate a new Elastic Beanstalk environment without Amazon RDS.

Answer: D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/decouple-rds-from-beanstalk/>

NEW QUESTION 268

- (Exam Topic 2)

A company is adding stored value for gift card) capability to its highly popular casual gaming website. Users need to be able to trade this value for other users' items on the platform. This would require both users' records be updated as a single transaction, or both users' records to be completely rolled back. Which AWS database options can provide the transactional capability required for this new feature? (Select TWO)

- A. Amazon DynamoDB with operations made with the ConsistentRead parameter set to true
- B. Amazon ElastiCache for Memcached with operations made within a transaction block
- C. Amazon Aurora MySQL with operations made within a transaction block
- D. Amazon DynamoDB with reads and writes made using Transact" operations
- E. Amazon Redshift with operations made within a transaction block

Answer: CD

NEW QUESTION 271

- (Exam Topic 2)

According to best practice, how should access keys be managed in AWS? (Choose two.)

- A. Use the same access key in all applications for consistency.
- B. Delete all access keys for the account root user.
- C. Leave unused access keys in the account for tracking purposes.
- D. Embed and encrypt access keys in code for continuous deployment.
- E. Use Amazon IAM roles instead of access keys where possible.

Answer: BE

NEW QUESTION 274

- (Exam Topic 2)

A front-end web application is using Amazon Cognito user pools to handle the user authentication flow. A developer is integrating Amazon DynamoDB into the application using the AWS SDK for JavaScript

How would the developer securely call the API without exposing the access or secret keys?

- A. Configure Amazon Cognito identity pools and exchange the JSON Web Token (JWT) for temporary credentials
- B. Run the web application in an Amazon EC2 instance with the instance profile configured
- C. Hardcode the credentials use Amazon S3 to host the web application, and enable server-side encryption
- D. Use Amazon Cognito user pool JSON Web Tokens (JWTs) to access the DynamoDB APIs.

Answer: C

NEW QUESTION 278

- (Exam Topic 2)

A Development team wants to instrument their code to provide more detailed information to AWS X-Ray than simple outgoing and incoming requests. This will generate large amounts of data, so the Development team wants to implement indexing so they can filter the data.

What should the Development team do to achieve this?

- A. Add annotations to the segment document and the code
- B. Add metadata to the segment document and the code
- C. Configure the necessary X-Ray environment variables
- D. Install required plugins for the appropriate AWS SDK

Answer: A

Explanation:

<https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-python-segment.html> <https://docs.aws.amazon.com/xray/latest/devguide/xray-concepts.html#xray-concepts-annotations>

NEW QUESTION 279

- (Exam Topic 2)

An organization is storing large files in Amazon S3, and is writing a web application to display meta-data about the files to end-users. Based on the metadata a user selects an object to download. The organization needs a mechanism to index the files and provide single-digit millisecond latency retrieval for the metadata. What AWS service should be used to accomplish this?

- A. Amazon DynamoDB
- B. Amazon EC2
- C. AWS Lambda
- D. Amazon RDS

Answer: A

Explanation:

Amazon DynamoDB is a fast and flexible NoSQL database service for all applications that need consistent, single-digit millisecond latency at any scale. It is a fully managed database and supports both document and key-value data models. Its flexible data model and reliable performance make it a great fit for mobile, web, gaming, ad-tech, Internet of Things (IoT), and many other applications.

References:

NEW QUESTION 284

- (Exam Topic 2)

A developer is setting up Amazon API Gateway for their company's products. The API will be used by registered developers to query and update their environments. The company wants to limit the amount of requests end users can send for both cost and security reasons. Management wants to offer registered developers the option of buying larger packages that allow for more requests.

How can the developer accomplish this with the LEAST amount of overhead management?

- A. Enable throttling for the API Gateway stage
- B. Set a value for both the rate and burst capacity
- C. If a registered user chooses a larger package, create a stage for them, adjust the values, and share the new URL with them.
- D. Set up Amazon CloudWatch API logging in API Gateway. Create a filter based on the user and requestTime fields and create an alarm on this filter. Write an AWS Lambda function to analyze the values and requester information, and respond accordingly. Set up the function as the target for the alarm. If a registered user chooses a larger package, update the Lambda code with the values.
- E. Enable Amazon CloudWatch metrics for the API Gateway stage. Set up CloudWatch alarms based off the Count metric and the ApiName, Method, Resource, and Stage dimensions to alerts when request rates pass the threshold. Set the alarm action to Deny. If a registered user chooses a larger package, create a user-specific alarm and adjust the values.
- F. Set up a default usage plan, specify values for the rate and burst capacity, and associate it with a stage. If a registered user chooses a larger package, create a custom plan with the appropriate values and associate the plan with the user.

Answer: D

NEW QUESTION 288

- (Exam Topic 2)

An application is running on an EC2 instance. The Developer wants to store an application metric in Amazon CloudWatch.

What is the best practice for implementing this requirement?

- A. Use the PUT Object API call to send data to an S3 bucket
- B. Use an event notification to invoke a Lambda function to publish data to CloudWatch.
- C. Publish the metric data to an Amazon Kinesis Stream using a PutRecord API call

- D. Subscribe a Lambda function that publishes data to CloudWatch.
- E. Use the CloudWatch PutMetricData API call to submit a custom metric to CloudWatc
- F. Provide the required credentials to enable the API call.
- G. Use the CloudWatch PutMetricData API call to submit a custom metric to CloudWatc
- H. Launch theEC2 instance with the required IAM role to enable the API call.

Answer: D

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html

NEW QUESTION 289

- (Exam Topic 2)

A website's page load times are gradually increasing as more users access the system at the same time. Analysis indicates that a user profile is being loaded from a database in all the web pages being visited by each user and this is increasing the database load and the page load latency. To address this issue the Developer decides to cache the user profile data.

Which caching strategy will address this situation MOST efficiently?

- A. Create a new Amazon EC2 Instance and run a NoSQL database on i
- B. Cache the profile data within this database using the write-through caching strategy.
- C. Create an Amazon ElastiCache cluster to cache the user profile dat
- D. Use a cache-aside caching strategy.
- E. Use a dedicated Amazon RDS instance for caching profile dat
- F. Use a write-through caching strategy.
- G. Create an ElastiCache cluster to cache the user profile dat
- H. Use a write-through caching strategy.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/Strategies.html>

NEW QUESTION 290

- (Exam Topic 2)

A company provides APIs as a service and commits to a service level agreement (SLA) with all its users. To comply with each SLA, what should the company do?

- A. Enable throttling limits for each method in Amazon API Getaway.
- B. Create a usage plan for each user and request API keys to access the APIs.
- C. Enable API rate limiting in Amazon cognito for each user.
- D. Enable default throttling limits for each stage after deploying the APIs.

Answer: D

NEW QUESTION 291

- (Exam Topic 2)

An Amazon DynamoDB table uses a Global Secondary Index (GSI) to support read queries. The primary table is write-heavy, whereas the GSI is used for read operations. Looking at Amazon CloudWatch metrics, the Developer notices that write operations to the primary table are throttled frequently under heavy write activity. However, write capacity units to the primary table are available and not fully consumed.

Why is the table being throttled?

- A. The GSI write capacity units are underprovisioned
- B. There are not enough read capacity units on the primary table
- C. Amazon DynamoDB Streams is not enabled on the table
- D. A large write operation is being performed against another table

Answer: A

Explanation:

<https://stackoverflow.com/questions/39582752/do-global-secondary-index-gsi-in-dynamodb-impact-tables-provi> <https://medium.com/@synchrophoto/amazon-dynamodb-provisioning-write-capacity-for-global-secondary-inde>

NEW QUESTION 293

- (Exam Topic 2)

A software engineer developed an AWS Lambda function in Node.js to do some CPU-intensive data processing. With the default settings, the Lambda function takes about 5 minutes to complete. Which approach should a developer take to increase the speed of completion?"

- A. Instead of using Node j
- B. rewrite the Lambda function using Python
- C. Instead of packaging the libraries in the ZIP file with the function move them to a Lambda layer and use the layer with the function.
- D. Allocate the maximum available CPU units lo the function
- E. Increase the available memory to the function.

Answer: D

NEW QUESTION 294

- (Exam Topic 2)

An application deployed on AWS Elastic Beanstalk experiences increased error rates during deployments of new application versions, resulting in service

degradation for users. The Development team believes that this is because of the reduction in capacity during the deployment steps. The team would like to change the deployment policy configuration of the environment to an option that maintains full capacity during deployment while using the existing instances. Which deployment policy will meet these requirements while using the existing instances?

- A. All at once
- B. Rolling
- C. Rolling with additional batch
- D. Immutable

Answer: D

Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

NEW QUESTION 297

- (Exam Topic 2)

A developer is testing a Docker-based application that uses the AWS SDK to interact with Amazon DynamoDB In the local development environment, the application has used IAM access keys The application is now ready for deployment onto an ECS cluster. How should the application authenticate with AWS services in production?

- A. Configure an ECS task IAM role for the application to use
- B. Refactor the application to call AWS STS AssumeRole based on an instance role
- C. Configure AWS access key/secret access key environment variables with new credentials
- D. Configure the credentials file with a new access key/secret access key

Answer: A

NEW QUESTION 298

- (Exam Topic 2)

A developer wants to build an application that will allow new users to register and create new user accounts. The application must also allow users with social media accounts to log in using their social media credentials. Which AWS service or feature can be used to meet these requirements?

- A. AWS IAM
- B. Amazon Cognito identity pools
- C. Amazon Cognito user pools
- D. AWS Directory Service

Answer: C

Explanation:

Reference:

<https://aws.amazon.com/blogs/apn/how-to-authenticate-users-into-your-apps-using-applicationload-balancer-and>

NEW QUESTION 302

- (Exam Topic 2)

A development team is working on a mobile app that allows users to upload pictures to Amazon S3 The team expects the app will be used by hundreds of thousands of users during a single event simultaneously Once the pictures are uploaded, the backend service will scan and parse the pictures for inappropriate content.

Which approach is the MOST resilient way to achieve this goal which also smooths out temporary volume spikes for the backend service?

- A. Develop an AWS Lambda function to check the upload folder in the S3 bucket
- B. If new uploaded pictures are detected, the Lambda function will scan and parse them
- C. Once a picture is uploaded to Amazon S3: publish the event to an Amazon SQS queue
- D. Use the queue as an event source to trigger an AWS Lambda function In the Lambda function scan and parse the picture.
- E. When the user uploads a picture, invoke an API hosted in Amazon API Gateway
- F. The API will invoke an AWS Lambda function to scan and parse the picture
- G. Create a state machine in AWS Step Functions to check the upload folder in the S3 bucket
- H. If a new picture is detected, invoke an AWS Lambda function to scan and parse it.

Answer: B

NEW QUESTION 306

- (Exam Topic 2)

An application needs to use the IP address of the client in its processing. The application has been moved into AWS and has been placed behind an Application Load Balancer (ALB). However, all the client IP addresses now appear to be the same. The application must maintain the ability to scale horizontally. Based on this scenario, what is the MOST cost-effective solution to this problem?

- A. Remove the application from the AL
- B. Delete the ALB and change Amazon Route 53 to direct traffic to the instance running the application.
- C. Remove the application from the AL
- D. Create a Classic Load Balancer in its place
- E. Direct traffic to the application using the HTTP protocol.
- F. Alter the application code to inspect the X-Forwarded-For header
- G. Ensure that the code can work properly if a list of IP addresses is passed in the header.
- H. Alter the application code to inspect a custom header
- I. Alter the client code to pass the IP address in the custom header.

Answer: C

NEW QUESTION 307

- (Exam Topic 2)

A Developer is writing a REST service that will add items to a shopping list. The service is built on Amazon API Gateway with AWS Lambda integrations. The shopping list items are sent as query string parameters in the method request.

How should the Developer convert the query string parameters to arguments for the Lambda function?

- A. Enable request validation
- B. Include the Amazon Resource Name (ARN) of the Lambda function
- C. Change the integration type
- D. Create a mapping template

Answer: D

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/integrating-api-with-aws-services-lambda.html#>

NEW QUESTION 311

- (Exam Topic 2)

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: C

Explanation:

Reference: <https://docs.aws.amazon.com/aws-technical-content/latest/microservices-on-aws/microservices-on-aw>

NEW QUESTION 314

- (Exam Topic 2)

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: A

NEW QUESTION 316

- (Exam Topic 2)

A developer is writing an application that will process data delivered into an Amazon S3 bucket. The data is delivered approximately 10 times a day, and the developer expects the data will be processed in less than 1 minute, on average.

How can the developer deploy and invoke the application with the lowest cost and lowest latency?

- A. Deploy the application as an AWS Lambda function and invoke it with an Amazon CloudWatch alarm triggered by an S3 object upload
- B. Deploy the application as an AWS Lambda function and invoke it with an S3 event notification
- C. Deploy the application as an AWS Lambda function and invoke it with an Amazon CloudWatch scheduled event
- D. Deploy the application onto an Amazon EC2 instance and have it poll the S3 bucket for new objects.

Answer: A

Explanation:

Reference: <https://docs.aws.amazon.com/lambda/latest/dg/with-s3.html>

NEW QUESTION 319

- (Exam Topic 2)

A team of Developers must migrate an application running inside an AWS Elastic Beanstalk environment from a Classic Load Balancer to an Application Load Balancer.

Which steps should be taken to accomplish the task using the AWS Management Console?

- A. *1. Update the application code in the existing deployment.* 2. Select a new load balancer type before running the deployment.* 3. Deploy the new version of the application code to the environment.
- B. *1. Create a new environment with the same configurations except for the load balancer type.* 2. Deploy the same application version as used in the original environment.* 3. Run the swap-environment-cnames action.
- C. *1. Clone the existing environment, changing the associated load balancer type.*2. Deploy the same application version as used in the original environment.*3. Run the swap-environment-cnames action.
- D. *1. Edit the environment definitions in the existing deployment.*2. Change the associated load balancer type according to the requirements.*3. Rebuild the environment with the new load balancer type.

Answer: B

Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.managing.elb.html>

By default, Elastic Beanstalk creates an Application Load Balancer for your environment when you enable load balancing with the Elastic Beanstalk console or the EB CLI. It configures the load balancer to listen for HTTP traffic on port 80 and forward this traffic to instances on the same port. You can choose the type of load balancer that your environment uses only during environment creation. Later, you can change settings to manage the behavior of your running environment's load balancer, but you can't change its type.

NEW QUESTION 324

- (Exam Topic 2)

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

Explanation:

Reference:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/install-CloudWatch-Agent-onpremise.Htm>

NEW QUESTION 325

- (Exam Topic 2)

A company needs to ingest terabytes of data each hour from thousands of sources that are delivered almost continually throughout the day. The volume of messages generated varies over the course of the day. Messages must be delivered in real time for fraud detection and live operational dashboards. Which approach will meet these requirements?

- A. Send the messages to an Amazon SQS queue, then process the messages by using a fleet of Amazon EC2 instances.
- B. Use the Amazon S3 API to write messages to an S3 bucket, then process the messages by using Amazon Redshift.
- C. Use AWS Data Pipeline to automate the movement and transformation of data.
- D. Use Amazon Kinesis Data Streams with Kinesis Client Library to ingest and deliver messages.

Answer: D

Explanation:

<https://aws.amazon.com/streaming-data/>

NEW QUESTION 330

- (Exam Topic 2)

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. Write a script that deletes old records; schedule the scripts as a cron job on an Amazon EC2 instance.
- B. Add an attribute with the expiration time; enable the Time To Live feature based on that attribute.
- C. Each day, create a new table to hold session data; delete the previous day's table.
- D. Add an attribute with the expiration time; name the attribute ItemExpiration.

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/time-to-live-ttl-how-to.html>

NEW QUESTION 332

- (Exam Topic 2)

An advertising company has a dynamic website with heavy traffic. The company wants to migrate the website infrastructure to AWS to handle everything except website development.

Which solution BEST meets these requirements?

- A. Use AWS VM Import to migrate a web server image to AWS. Launch the image on a compute-optimized Amazon EC2 instance.
- B. Launch multiple Amazon Lightsail instances behind a load balancer.
- C. Set up the website on those instances.
- D. Deploy the website code in an AWS Elastic Beanstalk environment.
- E. Use Auto Scaling to scale the number of instances.
- F. Use Amazon S3 to host the website.
- G. Use Amazon CloudFront to deliver the content at scale.

Answer: C

NEW QUESTION 336

- (Exam Topic 2)

A developer is refactoring a monolithic application. The application takes a POST request and performs several operations. Some of the operations are in parallel while others run sequentially. These operations have been refactored into individual AWS Lambda functions. The POST request will be processed by Amazon API Gateway.

How should the developer invoke the Lambda functions in the same sequence using API Gateway*?

- A. Use Amazon SQS to invoke the Lambda functions
- B. Use an AWS Step Functions activity to run the Lambda functions
- C. Use Amazon SNS to trigger the Lambda functions
- D. Use an AWS Step Functions state machine to orchestrate the Lambda functions.

Answer: A

NEW QUESTION 338

- (Exam Topic 2)

A company is developing an application that will be accessed through the Amazon API Gateway REST API. Registered users should be the only ones who can access certain resources of this API. The token being used should expire automatically and needs to be refreshed periodically. How can a developer meet these requirements?

- A. Create an Amazon Cognito identity pool, configure the Amazon Cognito Authorizer in API Gateway, and use the temporary credentials generated by the identity pool
- B. Create and maintain a database record for each user with a corresponding token and use an AWS Lambda authorizer in API Gateway
- C. Create an Amazon Cognito user pool, configure the Cognito Authorizer in API Gateway, and use the identity or access token
- D. Create an IAM user for each API user, attach an invoke permissions policy to the AP
- E. and use an IAM authorizer in API Gateway.

Answer: C

Explanation:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/cognito-custom-scopes-api-gateway/>

NEW QUESTION 340

- (Exam Topic 2)

A developer has written an Amazon Kinesis Data Streams application. As usage grows and traffic over time, the application is regularly receiving ProvisionedThroughputExceededException error messages.

Which steps should the Developer take to resolve the error? (Select Two.)

- A. Use Auto scaling to scale the stream for better performance.
- B. Increase the delay between the GetRecords call and the PutRecords call.
- C. Increase the number of shards in the data stream.
- D. Specify a shard iterator using the shardIterator parameter.
- E. Implement exponential backoff on the GetRecords call and the PutRecords call.

Answer: BD

Explanation:

Reference: <https://docs.aws.amazon.com/streams/latest/dev/troubleshooting-consumers.html>

NEW QUESTION 341

- (Exam Topic 2)

A developer has discovered that an application responsible for processing messages in an Amazon SQS queue is routinely falling behind. The application is capable of processing multiple messages in one execution, but is only receiving one message at a time.

What should the developer do to increase the number of messages the application receives?

- A. Call the ChangeMessageVisibility API for the queue and set MaxNumberOfMessages to a value greater than the default of 1.
- B. Call the AddPermission API to set MaxNumberOfMessages for the ReceiveMessage action to a value greater than the default of 1.
- C. Call the ReceiveMessage API to set MaxNumberOfMessages to a value greater than the default of 1
- D. Call the SetQueueAttributes API for the queue and set MaxNumberOfMessages to a value greater than the default of 1.

Answer: A

NEW QUESTION 344

- (Exam Topic 2)

A company is launching an ecommerce website and will host the static data in Amazon S3. The company expects approximately 1 000 transactions per second (TPS) for GET and PUT requests in total. Logging must be enabled to track all requests and must be retained for auditing purposes.

What is the MOST cost-effective solution?

- A. Enable AWS CloudTrail logging for the S3 bucket-level action and create a lifecycle policy to move the data from the log bucket to Amazon S3 Glacier in 90 days
- B. Enable S3 server access logging and create a lifecycle policy to expire the data in 90 days
- C. Enable AWS CloudTrail logging for the S3 bucket-level action and create a lifecycle policy to expire the data in 90 days
- D. Enable S3 server access logging and create a lifecycle policy to move the data to Amazon S3 Glacier in 90 days.

Answer: C

Explanation:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/dev/cloudtrail-request-identification.html>

NEW QUESTION 349

- (Exam Topic 2)

When developing an AWS Lambda function that processes Amazon Kinesis Data Streams, Administrators within the company must receive a notice that includes the processed data.

How should the Developer write the function to send processed data to the Administrators?

- A. Separate the Lambda handler from the core logic
- B. Use Amazon CloudWatch Events to send the processed data
- C. Publish the processed data to an Amazon SNS topic
- D. Push the processed data to Amazon SQS

Answer: C

Explanation:

<https://stackoverflow.com/questions/13681213/what-is-the-difference-between-amazon-sns-and-amazon-sqs> <https://stackoverflow.com/questions/31484868/can-you-publish-a-message-to-an-sns-topic-using-an-aws-lambda>

NEW QUESTION 350

- (Exam Topic 2)

A Developer is trying to make API calls using SDK. The IAM user credentials used by the application require multi-factor authentication for all API calls. Which method the Developer use to access the multi-factor authentication protected API?

- A. GetFederationToken
- B. GetCallerIdentity
- C. GetSessionToken
- D. DecodeAutherizationMessage

Answer: B

NEW QUESTION 352

- (Exam Topic 2)

The development team is working on an API that will be served from Amazon API gateway. The API will be served from three environments: development, test, and production. The API Gateway is configured to use 237 GB of cache in all three stages. Which is the MOST cost-efficient deployment strategy?

- A. Create a single API Gateway with all three stages.
- B. Create three API Gateways, one for each stage in a single AWS account.
- C. Create an API Gateway in three separate AWS accounts.
- D. Enable the cache for development and test environments only when needed.

Answer: D

NEW QUESTION 354

- (Exam Topic 2)

An on-premises legacy application is caching data files locally and writing shared images to local disks. What is necessary to allow for horizontal scaling when migrating the application to AWS?

- A. Modify the application to have both shared images and caching data written to Amazon EBS.
- B. Modify the application to read and write cache data on Amazon S3, and also store shared images on S3.
- C. Modify the application to use Amazon S3 for serving shared images; cache data can then be written to local disks.
- D. Modify the application to read and write cache data on Amazon S3, while continuing to write shared images to local disks.

Answer: C

NEW QUESTION 358

- (Exam Topic 2)

A development team wants to immediately build and deploy an application whenever there is a change to the source code. Which approaches could be used to trigger the deployment? (Select TWO.)

- A. Store the source code in an Amazon S3 bucket Configure AWS CodePipeline to start whenever a file in the bucket changes
- B. Store the source code in an encrypted Amazon EBS volume Configure AWS CodePipeline to start whenever a file in the volume changes
- C. Store the source code in an AWS CodeCommit repository Configure AWS CodePipeline to start whenever a change is committed to the repository.
- D. Store the source code in an Amazon S3 bucket Configure AWS CodePipeline to start every 15 minutes
- E. Store the source code in an Amazon EC2 instance's ephemeral storag
- F. Configure the instance to start AWS CodePipeline whenever there are changes to the source code

Answer: BC

Explanation:

Reference: <https://docs.aws.amazon.com/codepipeline/latest/userguide/tutorials-ecs-ecr-codedeploy.html>

NEW QUESTION 363

- (Exam Topic 2)

A company is running an application built on AWS Lambda functions. One Lambda function has performance issues when it has to download a 50MB file from the Internet in every execution. This function is called multiple times a second. What solution would give the BEST performance increase?

- A. Cache the file in the /tmp directory
- B. Increase the Lambda maximum execution time
- C. Put an Elastic Load Balancer in front of the Lambda function
- D. Cache the file in Amazon S3

Answer: A

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/runtimes-context.html>

NEW QUESTION 368

- (Exam Topic 2)

A company is using continuous integration and continuous delivery systems. A Developer now needs to automate a software package deployment to both Amazon EC2 instances and virtual servers running on-premises.

Which AWS service should be used to accomplish this?

- A. AWS CodePipeline
- B. AWS CodeBuild
- C. AWS Elastic Beanstalk
- D. AWS CodeDeploy

Answer: D

NEW QUESTION 370

- (Exam Topic 2)

A company has a REST application comprised of an Amazon API Gateway and several AWS Lambda functions. A developer is responding to an alert that the API Gateway's HTTP response error rate has unexpectedly increased. The developer must determine which Lambda function is malfunctioning.

Which method would help the developer make this determination while minimizing delays?

- A. Execute an Amazon Athena query against the API Gateway and Lambda execution logs.
- B. Execute an Amazon CloudWatch Logs Insights query against the API Gateway and Lambda execution logs.
- C. Download the API Gateway and Lambda execution logs from Amazon S3, and perform a line-by-line search against them.
- D. Download the API Gateway and Lambda execution logs from Amazon CloudWatch Events, and perform line-by-line search against them.

Answer: D

NEW QUESTION 372

- (Exam Topic 2)

An application is using Amazon DynamoDB as its data store, and should be able to read 100 items per second as strongly consistent reads. Each item is 5 KB in size.

To what value should the table's provisioned read throughput be set?

- A. 50 read capacity units
- B. 100 read capacity units
- C. 200 read capacity units
- D. 500 read capacity units

Answer: C

NEW QUESTION 376

- (Exam Topic 2)

A Developer has a stateful web server on-premises that is being migrated to AWS. The Developer must have greater elasticity in the new design.

How should the Developer re-factor the application to make it more elastic? (Choose two.)

- A. Use pessimistic concurrency on Amazon DynamoDB
- B. Use Amazon CloudFront with an Auto Scaling group
- C. Use Amazon CloudFront with an AWS Web Application Firewall
- D. Store session state data in an Amazon DynamoDB table
- E. Use an ELB with an Auto Scaling group

Answer: DE

NEW QUESTION 378

- (Exam Topic 2)

An application is using single -node Amazon ElastiCache for Redis instance to improve read performance. Over time, demand for the application has increased exponentially, which has increased the load on the ElastiCache instance. It is critical that this cache layer handles the load and is resilient in case of node failures. What can the Developer do to address the load and resiliency requirements?

- A. Add a read replica instance.
- B. Migrate to a Memcached cluster.
- C. Migrate to an Amazon ElastiCache service cluster.
- D. Vertically scale the ElastiCache instance.

Answer: A

Explanation:

Reference: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.Redis.Groups.html>

NEW QUESTION 380

- (Exam Topic 2)

A developer has created a new AWS IAM user that has s3 putobject permission to write to a specific Amazon bucket. This S3 bucket uses server-side encryption with AWS KMS managed keys (SEE-KMS) as the encryption. Using the access key and secret key of the IAM user, the application received an access denied error when calling the PutObject API.

How can this issue be resolved?

- A. Update the policy of the IAM user to allow the s3 Encrypt action.
- B. Update the bucket policy of the S3 bucket to allow the IAM user to upload objects
- C. Update the policy of the IAM user to allow the kms GenerateDatakey action
- D. Update the ACL of the bucket to allow the IAM user to upload objects

Answer: C

NEW QUESTION 382

- (Exam Topic 2)

A Developer decides to store highly secure data in Amazon S3 and wants to implement server-side encryption (SSE) with granular control of who can access the master key. Company policy requires that the master key be created, rotated, and disabled easily when needed, all for security reasons. Which solution should be used to meet these requirements?

- A. SSE with Amazon S3 managed keys (SSE-S3)
- B. SSE with AWS KMS managed keys (SSE KMS)
- C. SSE with AWS Secrets Manager
- D. SSE with customer provided encryption keys

Answer: B

NEW QUESTION 383

- (Exam Topic 2)

A company has a legacy application that was migrated to a fleet of Amazon EC2 instances. The application stores data in a MySQL database that is currently installed on a single EC2 instance. The company has decided to migrate the database from the EC2 instance to MySQL on Amazon RDS. What should the developer do to update the application to support data storage in Amazon RDS?

- A. Update the database connection parameters in the application to point to the new RDS instance
- B. Add a script to the EC2 instance that implements an AWS SDK for requesting database credentials.
- C. Create a new EC2 instance with an IAM role that allows access to the new RDS database
- D. Create an AWS Lambda function that will route traffic from the EC2 instance to the RDS database.

Answer: A

NEW QUESTION 385

- (Exam Topic 2)

After installing the AWS CLI, a Developer tries to run the command `aws configure` but receives the following error:

Error: aws: command not found

What is the most likely cause of this error?

- A. The aws executable is not in the PATH environment variable.
- B. Access to the aws executable has been denied to the installer.
- C. Incorrect AWS credentials were provided.
- D. The aws script does not have an executable file mode.

Answer: A

Explanation:

<https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-troubleshooting.html>

NEW QUESTION 387

- (Exam Topic 2)

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:

<https://aws.amazon.com/blogs/networking-and-content-delivery/amazon-s3-amazon-cloudfront-a-match-made-i>

NEW QUESTION 390

- (Exam Topic 2)

A Developer is creating a Lambda function that will generate and export a file. The function requires 100 MB of temporary storage for temporary files while executing. These files will not be needed after the function is complete.

How can the Developer MOST efficiently handle the temporary files?

- A. Store the files in EBS and delete the files at the end of the Lambda function.
- B. Copy the files to EFS and delete the files at the end of the Lambda function.
- C. Store the files in the /tmp directory and delete the files at the end of the Lambda function.
- D. Copy the files to an S3 bucket with a lifecycle policy to delete the files.

Answer: C

NEW QUESTION 394

- (Exam Topic 2)

An ecommerce startup is preparing for an annual sales event. As the traffic to the company's application increases, the development team wants to be notified when the Amazon EC2 instance's CPU utilization exceeds 80%.

Which solution will meet this requirement?

- A. Create a custom Amazon CloudWatch alarm that sends a notification to an Amazon SNS topic when the CPU utilization exceeds 80%.
- B. Create a custom AWS CloudTrail alarm that sends a notification to an Amazon SNS topic when the CPU utilization exceeds 80%.
- C. Create a cron job on the EC2 instance that executes the `--describe-instance-information` command on the host instance every 15 minutes and sends the results to an Amazon SNS topic.
- D. Create an AWS Lambda function that queries the AWS CloudTrail logs for the CPU utilization metric every 15 minutes and sends a notification to an Amazon SNS topic when the CPU utilization exceeds 80%.

Answer: C

NEW QUESTION 396

- (Exam Topic 2)

The upload of a 15 GB object to Amazon S3 fails. The error message reads: "Your proposed upload exceeds the maximum allowed object size."

What technique will allow the Developer to upload this object?

- A. Upload the object using the multi-part upload API.
- B. Upload the object over an AWS Direct Connect connection.
- C. Contact AWS Support to increase the object size limit.
- D. Upload the object to another AWS region.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/UploadingObjects.html>

NEW QUESTION 400

- (Exam Topic 2)

To include objects defined by the AWS Serverless Application Model (SAM) in an AWS CloudFormation template, in addition to Resources, what section MUST be included in the document root?

- A. Conditions
- B. Globals
- C. Transform
- D. Properties

Answer: C

Explanation:

<https://github.com/awslabs/serverless-application-model/blob/master/versions/2016-10-31.md> <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/sam-specification-template-an>

NEW QUESTION 405

- (Exam Topic 2)

A company requires that AWS Lambda functions written by developers log errors so system administrators can more effectively troubleshoot issues. What should the developers implement to meet this need?

- A. Publish errors to a dedicated Amazon SQS queue.
- B. Create an Amazon CloudWatch Events event to trigger based on certain Lambda events.
- C. Report errors through logging statements in Lambda function code.
- D. Set up an Amazon SNS topic that sends logging statements upon failure.

Answer: B

NEW QUESTION 410

- (Exam Topic 2)

A development team wants to run their container workloads on Amazon ECS. Each application container needs to share data with another container to collect logs and metrics.

What should the development team do to meet these requirements?

- A. Create two pod specifications. Make one to include the application container and the other to include the other container. Link the two pods together.
- B. Create two task definitions. Make one to include the application container and the other to include the other container.
- C. Mount a shared volume between the two tasks.
- D. Create one task definition. Specify both containers in the definition. Mount a shared volume between those two containers.
- E. Create a single pod specification. Include both containers in the specification. Mount a persistent volume to both containers.

Answer: C

NEW QUESTION 412

- (Exam Topic 2)

A developer is working on an AWS Lambda function that accesses Amazon DynamoDB. The Lambda function must retrieve an item and update some of its attributes, or create the item if it does not exist. The Lambda function has access to the primary key.

Which IAM permissions should the developer request for the Lambda function to achieve this functionality?

- A. dynamodb:DeleteItem dynamodb:GetItem dynamodb:PutItem
- B. dynamodb:UpdateItem dynamodb:GetItem dynamodb:DescribeTable
- C. dynamodb:GetRecords dynamodb:PutItem dynamodb:updateTable
- D. dynamodb:UpdateItem dynamodb:GetItem dynamodb:PutItem

Answer: C

Explanation:

Reference: <https://docs.aws.amazon.com/AWSJavaScriptSDK/latest/AWS/DynamoDB.html>

NEW QUESTION 417

- (Exam Topic 2)

A company is migrating from a monolithic architecture to a microservices-based architecture. The Developers need to refactor the application so that the many microservices can asynchronously communicate with each other without impacting performance. Use of which managed AWS services will enable asynchronous message passing? (Choose two.)

- A. Amazon SQS
- B. Amazon Cognito
- C. Amazon Kinesis
- D. Amazon SNS
- E. Amazon ElastiCache

Answer: AD

NEW QUESTION 418

- (Exam Topic 2)

An on-premises application makes repeated calls to store files to Amazon S3. As usage of the application has increased, "LimitExceeded" errors are being logged. What should be changed to fix this error?

- A. Implement exponential backoffs in the application.
- B. Load balance the application to multiple servers.
- C. Move the application to Amazon EC2.
- D. Add a one second delay to each API call.

Answer: A

NEW QUESTION 423

- (Exam Topic 2)

A developer is creating an AWS Lambda function that generates a new file each time it runs. Each new file must be checked into an AWS CodeCommit repository hosted in the same AWS account. How should the developer accomplish this?

- A. When the Lambda function starts, use the Git CLI to clone the repository
- B. Check the new file into the cloned repository and push the change.
- C. After the new file is created in Lambda, use cURL to invoke the CodeCommit API
- D. Send the file to the repository.
- E. Use an AWS SDK to instantiate a CodeCommit Client
- F. Invoke the put_file method to add the file to the repository.
- G. Upload the new file to an Amazon S3 bucket
- H. Create an AWS Step Function to accept S3 event
- I. In the Step Function, add the new file to the repository.

Answer: D

NEW QUESTION 424

- (Exam Topic 2)

A developer has built an application running on AWS Lambda using the AWS Serverless Application Model (AWS SAM). What is the correct order of execution to successfully deploy the application?

- A. * 1 Build the SAM template in Amazon EC2 * 2 Package the SAM template to Amazon EBS storage * 3. Deploy the SAM template from Amazon EBS.
- B. * 1 Build the SAM template locally * 2 Package the SAM template onto Amazon S3 * 3. Deploy the SAM template from Amazon S3.
- C. * 1 Build the SAM template locally * 2. Deploy the SAM template from Amazon S3. * 3 Package the SAM template for use
- D. * 1 Build the SAM template locally * 2 Package the SAM template from AWS CodeCommit
- E. * 3 Deploy the SAM template to CodeCommit

Answer: B

Explanation:

Reference:

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverlessdeploying.html>

NEW QUESTION 429

- (Exam Topic 2)

A company wants to migrate an imaging service to Amazon EC2 while following security best practices. The images are sourced and read from a non-public Amazon S3 bucket. What should a developer do to meet these requirements?

- A. Create an IAM user with read-only permissions for the S3 bucket Temporarily store the user credentials in the Amazon EBS volume of the EC2 instance

- B. Create an IAM user with read-only permissions for the S3 bucket
- C. Temporarily store the user credentials in the user data of the EC2 instance.
- D. Create an EC2 service role with read-only permissions for the S3 bucket Attach the role to the EC2 instance
- E. Create an S3 service role with read-only permissions for the S3 bucket Attach the role to the EC2 instance

Answer: A

NEW QUESTION 431

- (Exam Topic 2)

A developer is storing sensitive data generated by an application in Amazon S3. The developer wants to encrypt the data at rest. A company policy requires an audit trail of when the master key was used and by whom.

Which encryption option will meet these requirements?

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

Answer: B

NEW QUESTION 432

- (Exam Topic 2)

Queries to an Amazon DynamoDB table are consuming a large amount of read capacity. The table has a significant number of large attributes. The application does not need all of the attribute data.

How can DynamoDB costs be minimized while maximizing application performance?

- A. Batch all the writes, and perform the write operations when no or few reads are being performed.
- B. Create a global secondary index with a minimum set of projected attributes.
- C. Implement exponential backoffs in the application.
- D. Load balance the reads to the table using an Application Load Balancer.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/APIReference/query-api-troubleshooting.html>

NEW QUESTION 434

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Developer-Associate Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-Developer-Associate Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-Developer-Associate/>

Money Back Guarantee

AWS-Certified-Developer-Associate Practice Exam Features:

- * AWS-Certified-Developer-Associate Questions and Answers Updated Frequently
- * AWS-Certified-Developer-Associate Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Certified-Developer-Associate Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AWS-Certified-Developer-Associate Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year