

Amazon

Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate



NEW QUESTION 1

- (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 2

- (Exam Topic 1)

A Developer is writing a Linux-based application to run on AWS Elastic Beanstalk. Application requirements state that the application must maintain full capacity during updates while minimizing cost.

Which type of Elastic Beanstalk deployment policy should the Developer specify for the environment?

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

Answer: D

Explanation:

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

NEW QUESTION 3

- (Exam Topic 1)

A Developer is writing a mobile application that allows users to view images from an S3 bucket. The users must be able to log in with their Amazon login, as well as Facebook® and/or Google® accounts.

How can the Developer provide this authentication functionality?

- A. Use Amazon Cognito with web identity federation.
- B. Use Amazon Cognito with SAML-based identity federation.
- C. Use AWS IAM Access/Secret keys in the application code to allow Get* on the S3 bucket.
- D. Use AWS STS AssumeRole in the application code and assume a role with Get* permissions on the S3 bucket.

Answer: A

Explanation:

Reference:

<http://jayendrapatil.com/tag/iam-role/>

<https://docs.aws.amazon.com/sdk-for-javascript/v2/developer-guide/loading-browser-credentials-federated-id.ht>

NEW QUESTION 4

- (Exam Topic 1)

After launching an instance that you intend to serve as a NAT (Network Address Translation) device in a public subnet you modify your route tables to have the NAT device be the target of internet bound traffic of your private subnet. When you try and make an outbound connection to the Internet from an instance in the private subnet, you are not successful.

Which of the following steps could resolve the issue?

- A. Attaching a second Elastic Network interface (ENI) to the NAT instance, and placing it in the private subnet
- B. Attaching a second Elastic Network Interface (ENI) to the instance in the private subnet, and placing it in the public subnet
- C. Disabling the Source/Destination Check attribute on the NAT instance
- D. Attaching an Elastic IP address to the instance in the private subnet

Answer: C

Explanation:

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Instance.html#NATInstance

NEW QUESTION 5

- (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 6

- (Exam Topic 1) Company

C is currently hosting their corporate site in an Amazon S3 bucket with Static Website Hosting enabled. Currently, when visitors go to <http://www.companyc.com> the index.html page is returned. Company C now would like a new page welcome.html to be returned when a visitor enters <http://www.companyc.com> in the browser.

Which of the following steps will allow Company C to meet this requirement? Choose 2 answers

- A. Upload an html page named welcome.html to their S3 bucket
- B. Create a welcome subfolder in their S3 bucket
- C. Set the Index Document property to welcome.html
- D. Move the index.html page to a welcome subfolder
- E. Set the Error Document property to welcome.html

Answer: AC

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/WebsiteHosting.html> <https://docs.aws.amazon.com/AmazonS3/latest/dev/HostingWebsiteOnS3Setup.html>

NEW QUESTION 7

- (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 CodeCommi
- 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 CodeCommi
- K. Create a CodePipeline to automate unit testin
- L. Use Amazon CloudWatch to trigger notification of failure events.

Answer: D

NEW QUESTION 8

- (Exam Topic 1)

A Developer must repeatedly and consistently deploy a serverless RESTful API on AWS. Which techniques will work? (Choose two.)

- A. Define a Swagger fil
- B. Use AWS Elastic Beanstalk to deploy the Swagger file.
- C. Define a Swagger fil
- D. Use AWS CodeDeploy to deploy the Swagger file.
- E. Deploy a SAM template with an inline Swagger definition.
- F. Define a Swagger fil

- G. Deploy a SAM template that references the Swagger file.
- H. Define an inline Swagger definition in a Lambda function.
- I. Invoke the Lambda function.

Answer: CD

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/02/aws-serverless-application-model-aws-sam-supports-inline-swagger-definitions/> <https://aws.amazon.com/about-aws/whats-new/2017/02/aws-serverless-application-model-aws-sam-supports-inline-swagger-definitions/>

NEW QUESTION 9

- (Exam Topic 1)

When using a large Scan operation in DynamoDB, what technique can be used to minimize the impact of a scan on a table's provisioned throughput?

- A. Set a smaller page size for the scan
- B. Use parallel scans
- C. Define a range index on the table
- D. Prewarm the table by updating all items

Answer: A

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-query-scan.html>

Because a Scan operation reads an entire page (by default, 1 MB), you can reduce the impact of the scan operation by setting a smaller page size. The Scan operation provides a Limit parameter that you can use to set the page size for your request. Each Scan or Query request that has a smaller page size uses fewer read operations and creates a "pause" between each request. For example, if each item is 4 KB and you set the page size to 40 items, then a Query request would consume only 40 strongly consistent read operations or 20 eventually consistent read operations. A larger number of smaller Scan or Query operations would allow your other critical requests to succeed without throttling.

NEW QUESTION 10

- (Exam Topic 1)

A corporate web application is deployed within an Amazon VPC, and is connected to the corporate data center via IPSec VPN. The application must authenticate against the on-premise LDAP server. Once authenticated, logged-in users can only access an S3 keyspace specific to the user. Which two approaches can satisfy the objectives? Choose 2 answers

- A. The application authenticates against LDAP
- B. The application then calls the IAM Security Service to login to IAM using the LDAP credential
- C. The application can use the IAM temporary credentials to access the appropriate S3 bucket.
- D. The application authenticates against LDAP, and retrieves the name of an IAM role associated with the user
- E. The application then calls the IAM Security Token Service to assume that IAM Role
- F. The application can use the temporary credentials to access the appropriate S3 bucket.
- G. The application authenticates against IAM Security Token Service using the LDAP credential
- H. The application uses those temporary AWS security credentials to access the appropriate S3 bucket.
- I. Develop an identity broker which authenticates against LDAP, and then calls IAM Security Token Service to get IAM federated user credential
- J. The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket.
- K. Develop an identity broker which authenticates against IAM Security Token Service to assume an IAM Role to get temporary AWS security credential
- L. The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.

Answer: BD

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_request.html

NEW QUESTION 10

- (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 13

- (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 17

- (Exam Topic 1)

Where should the appspec.yml file be placed in order for AWS CodeDeploy to work?

- A. In the root of the application source code directory structure
- B. In the bin folder along with all the compiled code
- C. In an S3 bucket
- D. In the same folder as the application configuration files

Answer: A

NEW QUESTION 20

- (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 25

- (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 28

- (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 29

- (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 30

- (Exam Topic 1)

A Developer has setup an Amazon Kinesis Stream with 4 shards to ingest a maximum of 2500 records per second. A Lambda function has been configured to process these records.

In which order will these records be processed?

- A. Lambda will receive each record in the reverse order it was placed into the stream following a LIFO (last-in, first-out) method
- B. Lambda will receive each record in the exact order it was placed into the stream following a FIFO (first-in, first-out) method.
- C. Lambda will receive each record in the exact order it was placed into the shard following a FIFO (first-in, first-out) method
- D. There is no guarantee of order across shards.
- E. The Developer can select FIFO, (first-in, first-out), LIFO (last-in, last-out), random, or request specific record using the getRecords API.

Answer: C

NEW QUESTION 34

- (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 37

- (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 41

- (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. Amazon Device Messaging handle the registration.
- D. Implement a token vending service to handle the registration.
- E. 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 44

- (Exam Topic 1)

A Developer writes an AWS Lambda function and uploads the code in a .ZIP file to Amazon S3. The Developer makes changes to the code and uploads a new .ZIP file to Amazon S3. However, Lambda executes the earlier code.

How can the Developer fix this in the LEAST disruptive way?

- A. Create another Lambda function and specify the new .ZIP file.
- B. Call the update-function-code API.
- C. Remove the earlier .ZIP file first, then add the new .ZIP file.
- D. Call the create-alias API.

Answer: B

Explanation:

<https://docs.aws.amazon.com/cli/latest/reference/lambda/update-function-code.html>

NEW QUESTION 45

- (Exam Topic 1)

You are inserting 1000 new items every second in a DynamoDB table. Once an hour these items are analyzed and then are no longer needed. You need to minimize provisioned throughput, storage, and API calls.

Given these requirements, what is the most efficient way to manage these Items after the analysis?

- A. Retain the items in a single table
- B. Delete items individually over a 24 hour period
- C. Delete the table and create a new table per hour
- D. Create a new table per hour

Answer: C

NEW QUESTION 49

- (Exam Topic 1)

Company C has recently launched an online commerce site for bicycles on AWS. They have a "Product" DynamoDB table that stores details for each bicycle, such as, manufacturer, color, price, quantity and size to display in the online store. Due to customer demand, they want to include an image for each bicycle along with the existing details.

Which approach below provides the least impact to provisioned throughput on the "Product" table?

- A. Serialize the image and store it in multiple DynamoDB tables
- B. Create an "Images" DynamoDB table to store the Image with a foreign key constraint to the "Product" table
- C. Add an image data type to the "Product" table to store the images in binary format
- D. Store the images in Amazon S3 and add an S3 URL pointer to the "Product" table item for each image

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-use-s3-too.html#bp-use-s3-too-large-v> For example, consider the ProductCatalog table in the Creating Tables and Loading Data for Code Examples in DynamoDB section. Items in this table store information about item price, description, book authors, and dimensions for other products. If you wanted to store an image of each product that was too large to fit in an item, you could store the images in Amazon S3 instead of in DynamoDB.

NEW QUESTION 52

- (Exam Topic 1)

A company is developing a new online game that will run on top of Amazon ECS. Four distinct Amazon ECS services will be part of the architecture, each requiring specific permissions to various AWS services. The company wants to optimize the use of the underlying Amazon EC2 instances by bin packing the containers based on memory reservation.

Which configuration would allow the Development team to meet these requirements MOST securely?

- A. Create a new Identity and Access Management (IAM) instance profile containing the required permissions for the various ECS services, then associate that instance role with the underlying EC2 instances.
- B. Create four distinct IAM roles, each containing the required permissions for the associated ECS service, then configure each ECS service to reference the associated IAM role.
- C. Create four distinct IAM roles, each containing the required permissions for the associated ECS service, then, create an IAM group and configure the ECS cluster to reference that group.
- D. Create four distinct IAM roles, each containing the required permissions for the associated ECS service, then configure each ECS task definition to referene the associated IAM role.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-placement-strategies.html>.

NEW QUESTION 54

- (Exam Topic 1)

An application has hundreds of users. Each user may use multiple devices to access the application. The Developer wants to assign unique identifiers to these users regardless of the device they use.

Which of the following methods should be used to obtain unique identifiers?

- A. Create a user table in Amazon DynamoDB as key-value pairs of users and their device
- B. Use these keys as unique identifiers.
- C. Use IAM-generated access key IDs for the users as the unique identifier, but do not store secret keys.
- D. Implement developer-authenticated identities by using Amazon Cognito, and get credentials for these identities.
- E. Assign IAM users and roles to the user
- F. Use the unique IAM resource ID as the unique identifier.

Answer: C

NEW QUESTION 58

- (Exam Topic 1)

Which of the following are correct statements with policy evaluation logic in AWS Identity and Access Management? Choose 2 answers

- A. By default, all requests are denied
- B. An explicit allow overrides an explicit deny
- C. An explicit allow overrides default deny.
- D. An explicit deny does not override an explicit allow
- E. By default, all request are allowed

Answer: AC

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic.html

By default, all requests are implicitly denied. (Alternatively, by default, the AWS account root user has full access.) An explicit allow in an identity-based or resource-based policy overrides this default. If a permissions boundary, Organizations SCP, or session policy is present, it might override the allow with an implicit deny. An explicit deny in any policy overrides any allows.

NEW QUESTION 61

- (Exam Topic 1)

What is the format of structured notification messages sent by Amazon SNS?

- A. An XML object containing MessageId, UnsubscribeURL, Subject, Message and other values
- B. An JSON object containing MessageId, DuplicateFlag, Message and other values
- C. An XML object containing MessageId, DuplicateFlag, Message and other values
- D. An JSON object containing MessageId, unsubscribeURL, Subject, Message and other values

Answer: D

Explanation:

<https://docs.aws.amazon.com/sns/latest/dg/sns-message-and-json-formats.html#http-notification-json>

NEW QUESTION 65

- (Exam Topic 1)

A company has a multi-tiered web application on AWS. During a recent spike in traffic, one of the primary relational databases on Amazon RDS could not serve all the traffic. Some read queries for repeatedly accessed items failed, so users received error messages.

What can be done to minimize the impact on database read queries MOST efficiently during future traffic spikes?

- A. Use Amazon S3 to cache database query results.
- B. Use Amazon RDS as a custom origin for Amazon CloudFront.
- C. Use local storage and memory on Amazon EC2 instances to cache data.
- D. Use Amazon ElastiCache in front of the primary database to cache data.

Answer: D

NEW QUESTION 68

- (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 70

- (Exam Topic 1)

Which EC2 API call would you use to retrieve a list of Amazon Machine Images (AMIs)?

- A. DescribeInstances
- B. DescribeAMIs
- C. DescribeImages
- D. GetAMIs
- E. You cannot retrieve a list of AMIs as there are over 10,000 AMIs

Answer: C

Explanation:

https://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_DescribeImages.html

Describes the specified images (AMIs, AKIs, and ARIs) available to you or all of the images available to you.

NEW QUESTION 75

- (Exam Topic 1)

In a multi-container Docker environment in AWS Elastic Beanstalk, what is required to configure container instances in the environment?

- A. An Amazon ECS task definition
- B. An Amazon ECS cluster
- C. A Docker in an application package
- D. A CLI for Elastic Beanstalk

Answer: A

Explanation:

Reference: https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_ecs.html

NEW QUESTION 76

- (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 77

- (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>

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

NEW QUESTION 80

- (Exam Topic 1)

An AWS Lambda function must read data from an Amazon RDS MySQL database in a VPC and also reach a public endpoint over the internet to get additional data.

Which steps must be taken to allow the function to access both the RDS resource and the public endpoint? (Select TWO.)

- A. Modify the default configuration for the Lambda function to associate it with an Amazon VPC private subnet.
- B. Modify the default network access control list to allow outbound traffic.
- C. Add a NAT Gateway to the VPC.
- D. Modify the default configuration of the Lambda function to associate it with a VPC public subnet.
- E. Add an environmental variable to the Lambda function to allow outbound internet access.

Answer: AC

Explanation:

Reference: <https://docs.aws.amazon.com/lambda/latest/dg/vpc.html>

NEW QUESTION 85

- (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 89

- (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 91

- (Exam Topic 2)

A company is developing a report executed by AWS Step Functions Amazon CloudWatch shows errors in the Step Functions task state machine To troubleshoot each task, the state input needs to be included along with the error message in the state output.

Which coding practice can preserve both the original input and the error for the state?

- A. Use ResultPath in a Catch statement to include the error with the original input
- B. Use inputPath in a Catch statement and set the value to null.
- C. Use ErrorEquals in a Retry statement to include the error with the original input
- D. Use OutputPath in a Retry statement and set the value to \$.

Answer: A

Explanation:

Use ResultPath in a Catch to include the error with the original input.

Reference: <https://docs.aws.amazon.com/step-functions/latest/dg/input-output-resultpath.html>

NEW QUESTION 92

- (Exam Topic 2)

A company is using Amazon RDS MySQL instances for its application database tier and Apache Tomcat servers for its web tier. Most of the database queries from web applications are repeated read requests.

Use of which AWS service would increase in performance by adding in-memory store for repeated read queries?

- A. Amazon RDS Multi-AZ
- B. Amazon SQS
- C. Amazon ElastiCache
- D. Amazon RDS read replica

Answer: C

NEW QUESTION 96

- (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 98

- (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 102

- (Exam Topic 2)

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

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

The Developer needs to create/delete branches.

Which specific IAM permissions need to be added, based on the principle of least privilege?

- A. "codecommit:CreateBranch""codecommit>DeleteBranch"
- B. "codecommit:Put"
- C. "codecommit:Update"
- D. "codecommit:"

Answer: A

Explanation:

<https://docs.aws.amazon.com/codecommit/latest/userguide/auth-and-access-control-permissions-reference.html#>
<https://docs.aws.amazon.com/codecommit/latest/userguide/auth-and-access-control-iam-identity-based-access-c>

NEW QUESTION 106

- (Exam Topic 2)

An on-premises application is implemented using a Linux, Apache, MySQL and PHP (LAMP) stack. The Developer wants to run this application in AWS.

Which of the following sets of AWS services can be used to run this stack?

- A. Amazon API Gateway, Amazon S3
- B. AWS Lambda, Amazon DynamoDB
- C. Amazon EC2, Amazon Aurora
- D. Amazon Cognito, Amazon RDS
- E. Amazon ECS, Amazon EBS

Answer: C

NEW QUESTION 110

- (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 112

- (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 115

- (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 alert 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 119

- (Exam Topic 2)

A Developer is trying to deploy a serverless application using AWS CodeDeploy. The application was updated and needs to be redeployed.

What file does the Developer need to update to push that change through CodeDeploy?

- A. dockerrun.aws.json
- B. buildspec.yml
- C. appspec.yml
- D. ebextensions.config

Answer: C

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/application-revisions-push.html>

NEW QUESTION 124

- (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 CloudWatch
- F. Provide the required credentials to enable the API call.
- G. Use the CloudWatch PutMetricData API call to submit a custom metric to CloudWatch
- H. Launch the EC2 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 129

- (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 130

- (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 send 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 131

- (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 136

- (Exam Topic 2)

A company s website runs on an Amazon EC2 instance and uses Auto Scale the environment during peak times Website users across the world are experiencing high to sea latency due to static content on the EC2 instance, even during non-peak hours.

Which combination of steps will resolve the latency issue? (Select TWO)

- A. Double the Auto Scaling group's maximum number of servers.
- B. Host the application code on AWS Lambda
- C. Scale vertically by resizing the EC2 instances
- D. Create an Amazon CloudFront distribution to cache the static content
- E. Store the application's static content in Amazon S3

Answer: CD

Explanation:

Reference: <https://aws.amazon.com/getting-started/tutorials/deliver-content-faster/>

NEW QUESTION 138

- (Exam Topic 2)

An application is expected to process many files. Each file takes four minutes to process each AWS Lambda invocation. The Lambda function does not return any important data.

What is the fastest way to process all the files?

- A. First split the files to make them smaller, then process with synchronous RequestResponse Lambda invocations.
- B. Make synchronous RequestResponse Lambda invocations and process the files one by one.
- C. Make asynchronous Event Lambda invocations and process the files in parallel.
- D. First join all the files, then process it all at once with an asynchronous Event Lambda invocation.

Answer: C

NEW QUESTION 141

- (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 144

- (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 145

- (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 149

- (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 150

- (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 storage
- 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 152

- (Exam Topic 2)

A developer wants to send multi-value headers to an AWS Lambda function that is registered as a target with an Application Load Balancer (ALB). What should the developer do to achieve this?

- A. Place the Lambda function and target group in the same account
- B. Send the request body to the Lambda function with a size less than 1 MB
- C. Include the Base64 encoding status code, status description, and headers in the Lambda function
- D. Enable the multi-value headers on the ALB

Answer: D

NEW QUESTION 155

- (Exam Topic 2)

A company wants to implement authentication for its new REST service using Amazon API Gateway. To authenticate the calls, each request must include HTTP headers with a client ID and user ID. These credentials must be compared to authentication data in an Amazon DynamoDB table.

What MUST the company do to implement this authentication in API Gateway?

- A. Implement an AWS Lambda authorizer that references the DynamoDB authentication table
- B. Create a model that requires the credentials, then grant API Gateway access to the authentication table
- C. Modify the integration requests to require the credentials, then grant API Gateway access to the authentication table
- D. Implement an Amazon Cognito authorizer that references the DynamoDB authentication table

Answer: A

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-use-lambda-authorizer.html>

NEW QUESTION 156

- (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 161

- (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 166

- (Exam Topic 2)

A Developer must trigger an AWS Lambda function based on the item lifecycle activity in an Amazon DynamoDB table. How can the Developer create the solution?

- A. Enable a DynamoDB stream that publishes an Amazon SNS message
- B. Trigger the Lambda function synchronously from the SNS message.
- C. Enable a DynamoDB stream that publishes an SNS message
- D. Trigger the Lambda function asynchronously from the SNS message.
- E. Enable a DynamoDB stream, and trigger the Lambda function synchronously from the stream.
- F. Enable a DynamoDB stream, and trigger the Lambda function asynchronously from the stream.

Answer: C

Explanation:

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

NEW QUESTION 171

- (Exam Topic 2)

An application uses Amazon Kinesis Data Streams to ingest and process large streams of data records in real time. Amazon EC2 instances consume and process the data from the shards of the Kinesis data stream by using Amazon Kinesis Client Library (KCL). The application handles the failure scenarios and does not require standby workers. The application reports that a specific shard is receiving more data than expected. To adapt to the changes in the rate of data flow, the "hot" shard is resharded.

Assuming that the initial number of shards in the Kinesis data stream is 4, and after resharding the number of shards increased to 6, what is the maximum number of EC2 instances that can be deployed to process data from all the shards?

- A. 12
- B. 6
- C. 4
- D. 1

Answer: B

Explanation:

Typically, when you use the KCL, you should ensure that the number of instances does not exceed the number of shards (except for failure standby purposes). Each shard is processed by exactly one KCL worker and has exactly one corresponding record processor, so you never need multiple instances to process one shard. However, one worker can process any number of shards, so it's fine if the number of shards exceeds the number of instances.

<https://docs.aws.amazon.com/streams/latest/dev/kinesis-record-processor-scaling.html>

NEW QUESTION 173

- (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 176

- (Exam Topic 2)

An application is being developed to audit several AWS accounts. The application will run in Account A and must access AWS services in Accounts B and C. What is the MOST secure way to allow the application to call AWS services in each audited account?

- A. Configure cross-account roles in each audited account.
- B. Write code in Account A that assumes those roles.
- C. Use S3 cross-region replication to communicate among accounts, with Amazon S3 event notifications to trigger Lambda functions.
- D. Deploy an application in each audited account with its own role.
- E. Have Account A authenticate with the application.
- F. Create an IAM user with an access key in each audited account.
- G. Write code in Account A that uses those access keys.

Answer: A

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 178

- (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 180

- (Exam Topic 2)

An application uses Lambda functions to extract metadata from files uploaded to an S3 bucket; the metadata is stored in Amazon DynamoDB. The application starts behaving unexpectedly, and the Developer wants to examine the logs of the Lambda function code for errors.

Based on this system configuration, where would the Developer find the logs?

- A. Amazon S3
- B. AWS CloudTrail
- C. Amazon CloudWatch
- D. Amazon DynamoDB

Answer: C

NEW QUESTION 183

- (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 188

- (Exam Topic 2)

A development team is creating a new application designed to run on AWS. While the test and production environments will run on Amazon EC2 instances,

developers will each run their own environment on their laptops.

Which of the following is the simplest and MOST secure way to access AWS services from the local development machines?

- A. Use an IAM role to assume a role and execute API calls using the role.
- B. Create an IAM user to be shared with the entire development team, provide the development team with the access key.
- C. Create an IAM user for each developer on the team: provide each developer with a unique access key
- D. Set up a federation through an Amazon Cognito user pool.

Answer: A

NEW QUESTION 190

- (Exam Topic 2)

A developer wants the ability to roll back to a previous version of an AWS Lambda function in the event of errors caused by a new deployment.

How can the developer achieve this with MINIMAL impact on users?

- A. Change the application to use an alias that points to the current version Deploy the new version of the code Update the alias to use the newly deployed version
- B. If too many errors are encountered, point the alias back to the previous version
- C. Change the application to use an alias that points to the current version Deploy the new version of the code
- D. Update the alias to direct 10% of users to the newly deployed version
- E. If too many errors are encountered, send 100% of traffic to the previous version
- F. Do not make any changes to the application Deploy the new version of the code
- G. If too many errors are encountered, point the application back to the previous version using the version number in the Amazon Resource Name (ARN)
- H. Create three aliases: new, existing, and router Point the existing alias to the current version Have the router alias direct 100% of users to the existing alias Update the application to use the router alias Deploy the new version of the code Point the new alias to this version Update the router alias to direct 10% of users to the new alias If too many errors are encountered, send 100% of traffic to the existing alias

Answer: A

NEW QUESTION 193

- (Exam Topic 2)

A Developer is creating a template that uses AWS CloudFormation to deploy an application. This application is serverless and uses Amazon API Gateway, Amazon DynamoDB, and AWS Lambda.

Which tool should the Developer use to define simplified syntax for expressing serverless resources?

- A. CloudFormation serverless intrinsic functions
- B. AWS serverless express
- C. An AWS serverless application model
- D. A CloudFormation serverless plugin

Answer: A

NEW QUESTION 194

- (Exam Topic 2)

A developer has written an application that runs on Amazon EC2 instances and generates a value every minute. The Developer wants to monitor and graph the values generated over time without logging in to the instance each time.

Which approach should the Developer use to achieve this goal?

- A. Use the Amazon CloudWatch metrics reported by default for all EC2 instances View each value from the CloudWatch console.
- B. Develop the application to store each value in a file on Amazon S3 every minute with the Unix timestamp as the name
- C. Publish each generated value as a custom metric to Amazon CloudWatch using available AWS SDKs
- D. Store each value as a variable and add the variable to the list of EC2 metrics that should be reported to the Amazon CloudWatch console

Answer: C

NEW QUESTION 199

- (Exam Topic 2)

A company has an AWS CloudFormation template that is stored as a single file. The template is able to launch and create a full infrastructure stack.

Which best practice would increase the maintainability of the template?

- A. Use nested stacks for common template patterns.
- B. Embed credentials to prevent typos.
- C. Remove mappings to decrease the number of variables.
- D. Use AWS::Include to reference publicly-hosted template files.

Answer: A

NEW QUESTION 204

- (Exam Topic 3)

A developer has written the following IAM policy to provide access to an Amazon S3 bucket:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:PutObject"
      ],
      "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/secrets*"
    }
  ]
}
```

Which access does the policy allow regarding the s3:GetObject and s3:PutObject actions?

- A. Access on all buckets except the "DOC-EXAMPLE-BUCKET" bucket
- B. Access on all buckets that start with "DOC-EXAMPLE-BUCKET" except the "DOC-EXAMPLE-BUCKET/secrets" bucket
- C. Access on all objects in the "DOC-EXAMPLE-BUCKET" bucket along with access to all S3 actions for objects in the "DOC-EXAMPLE-BUCKET" bucket that start with "secrets"
- D. Access on all objects in the "DOC-EXAMPLE-BUCKET" bucket except on objects that start with "secrets"

Answer: D

Explanation:

Meaning:

DOC-EXAMPLE-BUCKET ==> bucket

DOC-EXAMPLE-BUCKET/* ==> contents in the bucket In this example,

ALLOW all "Objects" ==> DOC-EXAMPLE-BUCKET/*

DENY objects starting with secrets ==> DOC-EXAMPLE-BUCKET/secrets* <https://aws.amazon.com/blogs/security/iam-policies-and-bucket-policies-and-acls-oh-my-controlling-access-to-s>

NEW QUESTION 208

- (Exam Topic 3)

A developer is automating a new application deployment with AWS Serverless Application Model (AWS SAM) The new application has one AWS Lambda function and one Amazon S3 bucket The Lambda function must access the S3 bucket to only read objects How should the developer configure AWS SAM to grant the necessary read privilege to the S3 bucket?

- A. Reference a second Lambda authorizer function
- B. Add a custom S3 bucket policy to the Lambda function
- C. Create an Amazon Simple Queue Service (SQS) topic for only S3 object reads Reference the topic in the template
- D. Add the S3ReadPolicy template to the Lambda function's execution role

Answer: D

NEW QUESTION 213

- (Exam Topic 3)

A developer creates an Amazon S3 bucket to store project status files that are uploaded hourly. The developer also creates an AWS Lambda function that will be used to process the project status files

What should the developer do to invoke the function with the LEAST amount of AWS infrastructure'?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the function every 5 minutes and scan for new objects
- B. Create an S3 event notification to invoke the function when a new object is created in the S3 bucket
- C. Create an S3 event notification that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic Subscribe the function to the SNS topic.
- D. Create an S3 event notification that adds a message to an Amazon Simple Queue Service (Amazon SQS) queue Configure the function to poll the queue

Answer: B

NEW QUESTION 218

- (Exam Topic 3)

A developer is working on a serverless application that needs to process any changes to an Amazon DynamoDB table with an AWS Lambda function

How should the developer configure the Lambda function to detect changes to the DynamoDB table?

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

Answer: C

NEW QUESTION 222

- (Exam Topic 3)

A company is providing read access to objects in an Amazon S3 bucket for different customers. The company uses IAM permissions to restrict access to the S3 bucket. The customers can access only their own files.

Due to a regulation requirement, the company needs to enforce encryption in transit for interactions with Amazon S3. Which solution will meet these requirements?

- A. Add a bucket policy to the S3 bucket to deny S3 actions when the aws:SecureTransport condition is equal to false.
- B. Add a bucket policy to the S3 bucket to deny S3 actions when the s3:x-amz-aci condition is equal to public-read.
- C. Add an IAM policy to the IAM users to enforce the usage of the AWS SDK.
- D. Add an IAM policy to the IAM users that allows S3 actions when the s3:x-arnz-acl condition is equal to bucket-owner-read.

Answer: C

NEW QUESTION 226

- (Exam Topic 3)

A developer is deploying an application in the AWS Cloud by using AWS CloudFormation. The application will connect to an existing Amazon RDS database. The hostname of the RDS database is stored in AWS Systems Manager Parameter Store as a plaintext value. The developer needs to incorporate the database hostname into the CloudFormation template to initialize the application when the stack is created.

How should the developer reference the parameter that contains the database hostname?

- A. Use the ssm:dynamic reference.
- B. Use the Ref intrinsic function.
- C. Use the Fn:ImportValue intrinsic function.
- D. Use the ssm:secure dynamic reference.

Answer: C

NEW QUESTION 231

- (Exam Topic 3)

A developer is troubleshooting a three-tier application, which is deployed on Amazon EC2 instances. There is a connectivity problem between the application servers and the database servers.

Which AWS services or tools should be used to identify the faulty component? (Select TWO.)

- A. AWS CloudTrail.
- B. AWS Trusted Advisor.
- C. Amazon VPC Flow Logs.
- D. Network access control lists.
- E. AWS Config rules.

Answer: CD

NEW QUESTION 236

- (Exam Topic 3)

A physician's office management application requires that all data in transit between an EC2 instance and an Amazon EBS volume be encrypted.

Which of the following techniques fulfills this requirement? (Select TWO.)

- A. Create encrypted snapshots into Amazon S3.
- B. Use Amazon RDS with encryption.
- C. Use IAM roles to limit access to the Amazon EBS volume.
- D. Enable EBS encryption.
- E. Leverage OS-level encryption.

Answer: AD

NEW QUESTION 239

- (Exam Topic 3)

A company's e-commerce website is experiencing massive traffic spikes, which are causing performance problems in the company database. Users are reporting that accessing the website takes a long time.

A developer wants to implement a caching layer using Amazon ElastiCache. The website is required to be responsive no matter which product a user views, and the updates to product information and prices must be strongly consistent.

- A. Which cache writing policy will satisfy these requirements?
- B. Write to the cache directly and sync the backend at a later time.
- C. Write to the backend first and wait for the cache to expire.
- D. Write to the cache and the backend at the same time.
- E. Write to the backend first and invalidate the cache.

Answer: E

NEW QUESTION 242

- (Exam Topic 3)

A developer is writing an AWS Lambda function. The developer wants to log key events that occur during the Lambda function and include a unique identifier to associate the events with a specific function invocation.

Which of the following will help the developer accomplish this objective?

- A. Obtain the request identifier from the Lambda context object. Architect the application to write logs to the console.
- B. Obtain the request identifier from the Lambda event object. Architect the application to write logs to a file.
- C. Obtain the request identifier from the Lambda event object. Architect the application to write logs to the console.
- D. Obtain the request identifier from the Lambda context object. Architect the application to write logs to a file.

Answer: A

NEW QUESTION 243

- (Exam Topic 3)

A company has a three-tier application that is deployed in Amazon Elastic Container Service (Amazon ECS). The application is using an Amazon RDS for MySQL DB Instance. The application performs more database reads than writes.

During times of peak usage, the application's performance degrades. When this performance degradation occurs, the DB instance's ReadLatency metric in Amazon CloudWatch increases suddenly.

How should a developer modify the application to improve performance?

- A. Use Amazon ElastiCache to cache query results.
- B. Scale the ECS cluster to contain more ECS instances.
- C. Add read capacity units (RCUs) to the DB instance.
- D. Modify the ECS task definition to increase the task memory.

Answer: A

NEW QUESTION 245

- (Exam Topic 3)

A developer is designing a distributed application built using a microservices architect spanning multiple AWS accounts. The company's operations team wants to analyze and debug application issues from a centralized account.

How can the developer meet these requirements?

- A. Use an Amazon X-Ray agent with role assumption on to publish data into the centralized account.
- B. Use Amazon X-Ray and create a new IAM user to publish the access keys into the centralized account.
- C. Use VPC Flow Logs to collect application logs across different accounts.
- D. Enable AWS CloudTrail to publish the trails in an Amazon S3 bucket in the centralized account.

Answer: A

NEW QUESTION 246

- (Exam Topic 3)

A company has a web application that uses an Amazon Cognito user pool for authentication. The company wants to create a login page with the company logo. What should a developer do to meet these requirements?

- A. Create a hosted user interface in Amazon Cognito and customize it with the company logo.
- B. Create a login page with the company logo and upload it to Amazon Cognito.
- C. Create a login page in Amazon API Gateway with the logo and save the link in Amazon Cognito.
- D. Upload the logo to the Amazon Cognito app settings and point to the logo on a custom login page.

Answer: A

NEW QUESTION 250

- (Exam Topic 3)

A company has an application that is based on Amazon EC2. The company provides API access to the application through Amazon API Gateway and uses Amazon DynamoDB to store the application's data. A developer is investigating performance issues that are affecting the application. During peak usage, the application is overwhelmed by a large number of identical data read requests that come through APIs.

What is the MOST operationally efficient way for the developer to improve the application's performance?

- A. Use DynamoDB Accelerator (DAX) to cache database responses.
- B. Configure Amazon EC2 Auto Scaling policies to meet fluctuating demand.
- C. Enable API Gateway caching to cache API responses.
- D. Use Amazon ElastiCache to cache application responses.

Answer: D

NEW QUESTION 254

- (Exam Topic 3)

A developer is designing a web application in which new users will use their email addresses to create accounts. Millions of users are expected to sign up. The application will store attributes for each user.

Which AWS service or feature should the developer implement to meet these requirements?

- A. Amazon Cognito user pools
- B. AWS Mobile Hub User File Storage
- C. AWS AppSync
- D. AWS Mobile Hub Cloud Logic

Answer: A

NEW QUESTION 256

- (Exam Topic 3)

A company is developing a new web application in Python. A developer must deploy the application using AWS Elastic Beanstalk from the AWS Management Console. The developer creates an Elastic Beanstalk source bundle to upload using the console.

Which of the following are requirements when creating the source bundle? (Select TWO.)

- A. The source bundle must include the ebextensions.yaml file.
- B. The source bundle must not include a top-level directory.

- C. The source bundle must be compressed with any required dependencies in a top-level parent folder
- D. The source bundle must be created as a single zip or war file
- E. The source bundle must be uploaded into Amazon EFS.

Answer: BD

NEW QUESTION 260

- (Exam Topic 3)

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

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

A)

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

B)

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

C)

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

D)

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

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

Answer: A

NEW QUESTION 264

- (Exam Topic 3)

A company wants to migrate an existing web application to AWS. The application consists of two web servers and a MySQL database

The company wants the application to automatically scale in response to demand The company also wants to reduce its operational overhead for database backups and maintenance The company needs the ability to deploy multiple versions of the application concurrently

What is the MOST operationally efficient solution that meets these requirements?

- A. Deploy the application to AWS Elastic Beanstalk
- B. Migrate the database to an Amazon RDS Multi-AZ DB instance

- C. Create an Amazon Machine Image (AMI) that contains the application code
- D. Create an Auto Scaling group that is based on the AMI Integrate the Auto Scaling group with an Application Load Balancer for the web server
- E. Migrate the database to a MySQL instance that runs on an Amazon EC2 instance.
- F. Deploy the application to AWS Elastic Beanstalk
- G. Migrate the database to a MySQL instance that runs on an Amazon EC2 instance.
- H. Create an Amazon Machine Image (AMI) that contains the application code
- I. Create an Auto Scaling group that is based on the AMI
- J. Integrate the Auto Scaling group with an Application Load Balancer for the web server
- K. Migrate the database to an Amazon RDS Multi-AZ DB Instance

Answer: B

NEW QUESTION 265

- (Exam Topic 3)

How does Envelope Encryption work in AWS KMS?

- A. The Customer Master Key is used to encrypt/decrypt a data key The Plaintext Data Key is used to encrypt customer data.
- B. Two encryption keys are used The Customer Master Key encrypts customer data
- C. The Data Key is used to re-encrypt the encrypted data.
- D. Two encryption keys are used The Data Key encrypts customer data The Customer Master Key is used to re-encrypt the encrypted data
- E. The Customer Master Key is used to encrypt/decrypt a data key
- F. The Encrypted Data Key is used to encrypt customer data.

Answer: A

NEW QUESTION 267

- (Exam Topic 3)

A developer must cache dependent artifacts from Maven Central, a public package repository, as part of an application's build pipeline. The build pipeline has an AWS CodeArtifact repository where artifacts of the build are published. The developer needs a solution that requires minimum changes to the build pipeline Which solution meets these requirements?

- A. Modify the existing CodeArtifact repository to associate an upstream repository with the public package repository
- B. Create a new CodeArtifact repository that has an external connection to the public package repository
- C. Create a new CodeArtifact domain that contains a new repository that has an external connection to the public package repository
- D. Modify the CodeArtifact repository resource policy to allow artifacts to be fetched from the public package repository

Answer: D

NEW QUESTION 269

- (Exam Topic 3)

A developer is working on a serverless application. The application uses Amazon API Gateway, AWS Lambda functions that are written in Python, and Amazon DynamoDB.

Which combination of steps should the developer take so that the Lambda functions can be debugged in the event of application failures? (Select TWO)

- A. Configure an AWS CloudTrail trail to deliver log files to an Amazon S3 bucket
- B. Ensure that the Lambda functions write log messages to stdout and stderr
- C. Enable an AWS CloudTrail trail for the Lambda function
- D. Ensure that the execution role for the Lambda function has access to write to Amazon CloudWatch Logs.
- E. Use the Amazon CloudWatch metric for Lambda errors to create a CloudWatch alarm.

Answer: DE

NEW QUESTION 273

- (Exam Topic 3)

A developer has written an application hosted on Amazon EC2 instances. The application generates and uploads thousands of new objects to an Amazon S3 bucket located in the same AWS region. The size of each object is less than 1 MB. The application is taking too long to run. How can the performance of the application be improved?

- A. Use the S3 Multipart Upload API
- B. Use S3 Transfer Acceleration
- C. Upload the objects in parallel to Amazon S3
- D. Add a random prefix to the object keys

Answer: D

NEW QUESTION 275

- (Exam Topic 3)

A development team is migrating a monolithic application to Amazon API Gateway with AWS Lambda integrations using the AWS CDK The zip deployment package exceeds the Lambda direct upload deployment package size limit. How should the Lambda function be deployed?

- A. Use the zip file to create a Lambda layer and reference it using the `-code` CLI parameter
- B. Create a Docker image and reference the image using the `--docker-image` CLI parameter
- C. Upload a deployment package using the `--zip-file` CLI parameter
- D. Upload a deployment package to Amazon S3 and reference Amazon S3 using the `--code` CLI parameter

Answer: D

NEW QUESTION 280

- (Exam Topic 3)

A developer is designing an AWS Lambda function that create temporary files that are less than 10 MB during execution. The temporary files will be accessed and modified multiple times during execution. The developer has no need to save or retrieve these files in the future.

Where should the temporary file be stored?

- A. the /tmp directory
- B. Amazon EFS
- C. Amazon EBS
- D. Amazon S3

Answer: A

NEW QUESTION 282

- (Exam Topic 3)

An organization is using Amazon API Gateway to provide a public API called "Survey" for collecting user feedback posts about its products. The survey API has "DEV" and "PROD" stages and consists of one resource "/feedback" which allows users to retrieve/create/update single feedback posts.

A version-controlled Swagger file is used to define a new API that retrieves multiple feedback posts. To add the new API resource "/listFeedbackForProduct" the developer makes changes to the Swagger file, defines an API, uploads the file to the organization's version control system, and uses the API Gateway Import API feature to apply the changes to the Survey API. After successful import, the developer runs the tests against the DEV stage and finds that resource "/listFeedbackForProduct" is not available.

What is MOST likely the reason for resource not being available?

- A. Even though the Swagger import was successful, resource creation failed afterwards.
- B. There is a propagation delay of several minutes in creating API Gateway resources after import.
- C. The developer needs to restart the API Gateway stage after import in order to apply the changes.
- D. The developer needs to create a new deployment after import in order to deploy the changes.

Answer: A

NEW QUESTION 287

- (Exam Topic 3)

A developer tested an application locally and then deployed it to AWS Lambda. While testing the application remotely, the Lambda function fails with an access denied message. How can this issue be addressed?

- A. Update the Lambda function's execution role to include the missing permissions.
- B. Update the Lambda function's resource policy to include the missing permissions.
- C. Include an IAM policy document at the root of the deployment package and redeploy the Lambda function.
- D. Redeploy the Lambda function using an account with access to the AdministratorAccess policy.

Answer: A

NEW QUESTION 288

- (Exam Topic 3)

A company uses a third-party tool to build, bundle, and package its applications on-premises and store them locally. The company uses Amazon EC2 instances to run its front-end applications. How can an application be deployed from the source control system onto the EC2 instances?

- A. Use AWS CodeDeploy and point it to the local storage to directly deploy a bundle in a zip
- B. tar
- C. or tar.gz format
- D. Upload the bundle to an Amazon S3 bucket and specify the S3 location when doing a deployment using AWS CodeDeploy
- E. Create a repository using AWS CodeCommit to automatically trigger a deployment to the EC2 instances
- F. Use AWS CodeBuild to automatically deploy the latest build to the latest EC2 instances

Answer: B

NEW QUESTION 292

- (Exam Topic 3)

Multiple development teams are working on a project to migrate a monolithic application to a microservices-based application running on AWS Lambda. The teams need a way to centrally manage code that is shared across multiple functions.

Which approach requires the LEAST maintenance?

- A. Each team maintains the code for the common components in their own code repository.
- B. They build and deploy the components with their Lambda functions together.
- C. One team builds a Lambda layer to include the common components and shares the layer with the other teams.
- D. Each team builds and publishes the component they want to share to an Amazon S3 bucket. The Lambda functions will download the components from the bucket.
- E. One team builds a Docker container for the common components and shares the container with the other teams.

Answer: C

NEW QUESTION 295

- (Exam Topic 3)

A company has a large number of documents that are stored securely in Amazon S3. The company is creating an application that occasionally will read these documents. The application will be deployed on Amazon EC2 instances.

The company's security requirements mandate that no long-term credentials can be stored on the EC2 instances and that only the needed documents can be accessed. Only authorized users and applications can access the documents, access must be logged by Amazon S3, and each document must follow S3 Lifecycle policies for archival and destruction.

What should a developer do to meet these requirements?

- A. Create an event to invoke an AWS Lambda function when a document is uploaded Configure the function to write the documents to an Amazon Elastic File System (Amazon EFS) file system Configure the EC2 instances to mount the EFS file system Configure the application to access the documents that are stored in the file system as needed
- B. Create a user that has programmatic credentials, and attach a policy that allows read access to the S3 bucket Use the AWS CLI to configure those credentials for the EC2 instances to use Create an Amazon Machine Image (AMI), and add the access key and secret access key to the user data section to create environment variables Use the AMI to launch each EC2 instance that runs the application Add application code to use the keys that are stored in the environment variables to access the S3 bucket objects as needed.
- C. Modify the S3 bucket, make the bucket public, and make each object public Add application code to make REST calls to access the objects in the S3 bucket as needed
- D. Create an IAM role with permissions to read objects from Amazon S3 Attach the role to the EC2 instances as an instance profile Add application code to access the objects in the S3 bucket as needed.

Answer: D

NEW QUESTION 300

- (Exam Topic 3)

A developer is using Amazon S3 as the event source that invokes a Lambda function when new objects are created in the bucket The event source mapping information is stored in the bucket notification configuration The developer is working with different versions of the Lambda function, and has a constant need to update notification configuration so that Amazon S3 invokes the correct version

What is the MOST efficient and effective way to achieve mapping between the S3 event and Lambda?

- A. Use a different Lambda trigger
- B. Use Lambda environment variables
- C. Use a Lambda alias
- D. Use Lambda tags.

Answer: C

NEW QUESTION 305

- (Exam Topic 3)

A developer is building a backend system for the long-term storage of information from an inventory management system. The information needs to be stored so that other teams can build tools to report and analyze the data

How should the developer implement this solution to achieve the FASTEST running time?

- A. Create an AWS Lambda function that writes to Amazon S3 synchronously Increase the function's concurrency to match the highest expected value of concurrent scans and requests.
- B. Create an AWS Lambda function that writes to Amazon S3 asynchronously Configure a dead-letter queue to collect unsuccessful invocations
- C. Create an AWS Lambda function that writes to Amazon S3 synchronously Set the inventory system to retry failed requests.
- D. Create an AWS Lambda function that writes to an Amazon ElastiCache for Redis cluster asynchronously Configure a dead-letter queue to collect unsuccessful invocations.

Answer: A

NEW QUESTION 309

- (Exam Topic 3)

A developer is creating a role to access Amazon S3 buckets To create the role, the developer uses the AWS CLI create-role command. Which policy should be added to allow the Amazon EC2 service to assume the role?

- A. Managed policy
- B. Trust policy
- C. Inline policy
- D. Service control policy (SCP)

Answer: B

Explanation:

A JSON policy document in which you define the principals that you trust to assume the role. A role trust policy is a required resource-based policy that is attached to a role in IAM. The principals that you can specify in the trust policy include users, roles, accounts, and services.

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_terms-and-concepts.html

NEW QUESTION 310

- (Exam Topic 3)

A developer is creating a serverless web application and maintains different branches of code The developer wants to avoid updating the Amazon API Gateway target endpoint each time a new code push is performed What solution would allow the developer to perform a code push efficiently, without the need to update the API Gateway?

- A. Associate different AWS Lambda functions to an API Gateway target endpoint
- B. Create different stages in API Gateway, then associate API Gateway with AWS Lambda
- C. Create aliases and versions in AWS Lambda.
- D. Tag the AWS Lambda functions with different names

Answer: B

NEW QUESTION 311

- (Exam Topic 3)

An application is experiencing performance issues based on increased demand. This increased demand is on read-only historical records pulled from an Amazon

RDS-hosted database with custom views and queries. A developer improve performance without changing the database structure. Which approach will improve performance and MINIMIZE management overhead?

- A. Deploy Amazon DynamoDB, move all the data, and point to DynamoDB.
- B. Deploy Amazon ElasticCache for Redis and cache the data for the application.
- C. Deploy Memcached on Amazon EC2 and cache the data for the application.
- D. Deploy Amazon DynamoDB Accelerator (DAX) on Amazon RDS to improve cache performance

Answer: B

NEW QUESTION 315

- (Exam Topic 4)

A developer wants to migrate a Windows .NET application that is running on IIS with a Microsoft SQL Server database to AWS. The developer does not want to think about provisioning and managing the infrastructure.

What should the developer do to migrate the application with the LEAST amount of effort?

- A. Launch Amazon EC2 instances for Windows Serve
- B. Back up and restore the database to Amazon RD
- C. Deploy the web application to the new EC2 instances
- D. Back up and restore the database to Amazon RD
- E. Use the .NET Migration Assistant for AWS Elastic Beanstalk to migrate the web application to a preconfigured solution stack that Elastic Beanstalk provides.
- F. Migrate the database to Amazon DynamoDB Use Amazon API Gateway and AWS Lambda to create a web application interface that is hosted in an Amazon S3 bucket.
- G. Containerize the application on premise
- H. Push the image to Amazon Elastic Container Registry (Amazon ECR). Create an AWS CloudFormation template to deploy the application

Answer: B

NEW QUESTION 316

- (Exam Topic 4)

A security review for a software company's application infrastructure shows that there is no test coverage in any of the company's deployment pipelines. A developer must fix this issue as soon as possible. The company has been integrating the AWS Cloud Development Kit (AWS CDK) into the deployment process. However, much of the pipeline still uses AWS CloudFormation templates. The developer needs to add test coverage to all the deployment code.

Which solution will meet these requirements with the LEAST amount of configuration?

- A. Write unit tests by using the AWS CDK assertions modul
- B. Create CloudFormation template instances by using the module's Template class for the existing CloudFormation templates and the module's Capture class for the CDK stacks.
- C. Write unit tests by using the AWS CDK assertions modul
- D. Create CloudFormation template instances by using the module's Template.fromStringO method for the existing CloudFormation templates and the module's Template fromStackQ method for the CDK stacks.
- E. Convert the CloudFormation templates into CDK stacks by using the AWS CDK CfnInclude construct. Write unit tests against the templates by using CloudFormation rule assertions.
- F. Convert the CDK stacks into CloudFormation templates by using the AWS CDK CfnInclude construct Write unit tests against the templates by using CloudFormation rule assertions

Answer: A

NEW QUESTION 318

- (Exam Topic 4)

A developer is running an application on an Amazon EC2 instance. When the application tries to read an Amazon S3 bucket, the application fails. The developer notices that the associated IAM role is missing the S3 read permission. The developer needs to give the application the ability to read the S3 bucket.

Which solution will meet this requirement with the LEAST application disruption?

- A. Add the permission to the rol
- B. Terminate the existing EC2 instanc
- C. Launch a new EC2 instance.
- D. Add the permission to the role so that the change will take effect automatically.
- E. Add the permission to the rol
- F. Hibernate and restart the existing EC2 instance.
- G. Add the permission to the S3 bucke
- H. Restart the EC2 instance.

Answer: B

NEW QUESTION 323

- (Exam Topic 4)

A company has an online order website that uses Amazon DynamoDB to store item inventory. A sample of the inventory object is as follows:

```
{
  "Id": { "N": "456"},
  "Price": {"N": "650"},
  "ProductCategory": {"S": "Sporting Goods"}
}
```

A developer needs to reduce all inventory prices by 100 as long as the resulting price would not be less than 500. What should the developer do to make this change with the LEAST number of calls to DynamoDB?

- A. Perform a DynamoDB Query operation with the I
- B. If the price is ≥ 600 , perform an UpdateItem operation to update the price.
- C. Perform a DynamoDB UpdateItem operation with a condition expression of "Price ≥ 600 ".
- D. Perform a DynamoDB UpdateItem operation with a condition expression of "ProductCategory IN<{"S": "Sporting Goods"}> and Price = 600".

Answer: C

NEW QUESTION 324

- (Exam Topic 4)

A data-processing application includes an AWS Lambda function that processes data in several steps. Recently, the function has been reaching the Lambda timeout. A developer wants to use AWS X-Ray to find out how long each step is taking so that the developer can determine which step is causing the timeout. Which combination of actions should the developer take to accomplish this goal? (Select TWO.)

- A. Modify the application to call the PutMetricData API operation after each processing step
- B. Include the time taken in milliseconds.
- C. Use the aws lambda update-function-configuration AWS CLI command to enable active tracing on the Lambda function.
- D. Modify the application to record each processing step in an X-Ray subsegment by using the X-Ray software development kit (SDK).
- E. Add the xray:PutTraceSegments permission and the xray:PutTelemetryRecords permission to the Lambda function's execution role.
- F. Modify the application to put each processing step in a separate Lambda layer
- G. Include all the layers in the Lambda function.

Answer: B

NEW QUESTION 325

- (Exam Topic 4)

A company has a serverless application that uses Amazon API Gateway backed by AWS Lambda proxy integration. The company is developing several backend APIs. The company needs a landing page to provide an overview of navigation to the APIs.

A developer creates a new /LandingPage resource and a new GET method that uses mock integration. What should the developer do next to meet these requirements?

- A. Configure the integration request mapping template with Content-Type of text/html and statusCode of 200. Configure the integration response mapping template with Content-Type of application/json
- B. In the integration response mapping template, include the LandingPage HTML code that references the APIs.
- C. Configure the integration request mapping template with Content-Type of application/json
- D. In the integration request mapping template, include the LandingPage HTML code that references the API
- E. Configure the integration response mapping template with Content-Type of text/html and statusCode of 200.
- F. Configure the integration request mapping template with Content-Type of application/json and statusCode of 200. Configure the integration response mapping template with Content-Type of text/html
- G. In the integration response mapping template, include the LandingPage HTML code that references the APIs.
- H. Configure the integration request mapping template with Content-Type of text/html
- I. In the integration request mapping template, include the LandingPage HTML code that references the APIs. Configure the integration response mapping template with Content-Type of application/json and statusCode of 200.

Answer: B

NEW QUESTION 327

- (Exam Topic 4)

A company is using continuous integration/continuous delivery (CI/CD) systems. A that run on premises. Which AWS service should the developer use to meet these requirements?

- A. AWS Cloud9
- B. AWS CodeBuild
- C. AWS Elastic Beanstalk
- D. AWS CodeDeploy

Answer: D

NEW QUESTION 328

- (Exam Topic 4)

A company's security policies require all database passwords to be rotated every 30 days. The company uses different database platforms, including Amazon Aurora databases and proprietary NoSQL document databases, for different applications. A developer needs to implement a solution for password rotation. Which solution will meet these requirements?

- A. Create an AWS Lambda rotation function that has appropriate IAM permissions. Store the password in AWS Secrets Manager. Configure Secrets Manager to rotate the password by using the Lambda function.
- B. Encrypt the existing password with AWS Key Management Service (AWS KMS). Export the existing password. Generate a random password with AWS KMS. Use the AWS KMS password renewal feature to replace the existing password with the new password.
- C. Create an AWS Lambda rotation function that has appropriate IAM permissions. Store the password in AWS Systems Manager Parameter Store. Configure Parameter Store to rotate the password by using the Lambda function.
- D. Integrate AWS Systems Manager Parameter Store with a Key Management Interoperability Protocol (KMIP)-compliant third-party secret manager to enable third-party database password rotation on AWS.

Answer: C

NEW QUESTION 329

- (Exam Topic 4)

A developer deploys a custom application to three Amazon EC2 instances. The application processes messages from an Amazon Simple Queue Service (Amazon SQS) standard queue with default settings. When the developer runs a load test on the Amazon SQS queue, the developer discovers that the application

processes many messages multiple times.
How can the developer ensure that the application processes each message exactly once?

- A. Modify the SQS standard queue to an SQS FIFO queue.
- B. Process the messages on one EC2 instance instead of three instances.
- C. Create a new SQS FIFO queue.
- D. Point the application to the new queue.
- E. Increase the DelaySeconds value on the current SQS queue.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues-moving.html> Moving from a standard queue to a FIFO queue:
If you have an existing application that uses standard queues and you want to take advantage of the ordering or exactly-once processing features of FIFO queues, you need to configure the queue and your application correctly.

Note:

You can't convert an existing standard queue into a FIFO queue. To make the move, you must either create a new FIFO queue for your application or delete your existing standard queue and recreate it as a FIFO queue.

NEW QUESTION 331

- (Exam Topic 4)

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

Which HTTP header should the developer use for this analysis?

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

Answer: C

Explanation:

<https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/X-Forwarded-Proto>

NEW QUESTION 333

- (Exam Topic 4)

A company has a front-end application that runs on four Amazon EC2 instances behind an Elastic Load Balancer (ELB) in a production environment that is provisioned by AWS Elastic Beanstalk. A developer needs to deploy and test new application code while updating the Elastic Beanstalk platform from the current version to a newer version of Node.js. The solution must result in zero downtime for the application.

Which solution meets these requirements?

- A. Clone the production environment to a different platform version.
- B. Deploy the new application code, and test it.
- C. Swap the environment URLs upon verification.
- D. Deploy the new application code in an all-at-once deployment to the existing EC2 instance.
- E. Test the code.
- F. Redeploy the previous code if verification fails.
- G. Perform an immutable update to deploy the new application code to new EC2 instances.
- H. Serve traffic to the new instances after they pass health checks.
- I. Use a rolling deployment for the new application code.
- J. Apply the code to a subset of EC2 instances until the tests pass.
- K. Redeploy the previous code if the tests fail.

Answer: D

NEW QUESTION 336

- (Exam Topic 4)

A developer is building a three-tier application with an Application Load Balancer (ALB), Amazon EC2 instances, and Amazon RDS. There is an alias record in Amazon Route 53 that points to the ALB. When the developer tries to access the ALB from a laptop, the request times out.

Which logs should the developer investigate to verify that the request is reaching the AWS network?

- A. VPC Flow Logs
- B. Amazon Route 53 logs
- C. AWS Systems Manager Agent logs
- D. Amazon CloudWatch agent logs

Answer: A

NEW QUESTION 337

- (Exam Topic 4)

A company is developing a microservice that will manage customer account data in an Amazon DynamoDB table. Insert, update, and delete requests will be rare. Read traffic will be heavy. The company must have the ability to access customer data quickly by using a customer ID. The microservice can tolerate stale data. Which solution will meet these requirements with the FEWEST possible read capacity units (RCUs)?

- A. Read the table by using eventually consistent reads.
- B. Read the table by using strongly consistent reads.
- C. Read the table by using transactional reads.
- D. Read the table by using strongly consistent PartiQL queries.

Answer: A

Explanation:

Key points: "Read heavy", "access data quickly", "can tolerate stale data" To achieve: "FEWEST" possible (RCUs)
For items up to 4 KB in size, one RCU can perform one strongly consistent read request per second. For items up to 4 KB in size, one RCU can perform two eventually consistent read requests per second. Transactional read requests require two RCUs to perform one read per second for items up to 4 KB. For example, a strongly consistent read of an 8 KB item would require two RCUs, an eventually consistent read of an 8 KB item would require one RCU, and a transactional read of an 8 KB item would require four RCUs. <https://aws.amazon.com/dynamodb/pricing/provisioned/>

NEW QUESTION 339

- (Exam Topic 4)

A developer is writing an AWS Lambda function. The Lambda function needs to access items that are stored in an Amazon DynamoDB table. What is the MOST secure way to configure this access for the Lambda function?

- A. Create an IAM user that has permissions to access the DynamoDB tabl
- B. Create an access key for this use
- C. Store the access key ID and secret..... key in the Lambda function environment variables.
- D. Add a resource-based policy to the DynamoDB table to allow access from the Lambda function's IAM role.
- E. Create an IAM policy that allows access to the DynamoDB tabl
- F. Attach this policy to the Lambda function's 1AM role.
- G. Create a DynamoDB Accelerator (DAX) cluste
- H. Configure the Lambda function to use the DAX cluster to access the DynamoDB table.

Answer: A

NEW QUESTION 344

- (Exam Topic 4)

A company manages a financial services application that stores a large volume of data in an Amazon DynamoDB table. A developer needs to improve the performance of the DynamoDB read queries without increasing the cost. Which solution meets these requirements?

- A. Use parallel scans
- B. Add a local secondary index (LSI).
- C. Create a DynamoDB Accelerator (DAX) cluster.
- D. Query with the Projection Expression parameter

Answer: C

NEW QUESTION 347

- (Exam Topic 4)

A developer deployed an application to an Amazon EC2 instance. The application needs to know the public IPv4 address of the instance. How can the application find this information?

- A. Query the instance metadata from <http://169.254.169.254/latest/meta-data/>.
- B. Query the instance user data from <http://169.254.169.254/latest/user-data/>.
- C. Query the Amazon Machine Image (AMI) information from <http://169.254.169.254/latest/meta-data/ami/>.
- D. Check the hosts file of the operating system.

Answer: A

NEW QUESTION 349

- (Exam Topic 4)

A company is using Amazon Cognito user pools for sign-up and login functionality for a web application. The company is using Amazon RDS for the application's data persistence and is using Amazon API Gateway and AWS Lambda for the application's API functionality. Users must provide their first name, last name, email address, and phone number to sign up. All API endpoints have a Cognito user pool authorizer to guard against unauthenticated requests.

A developer wants to show a personalized welcome screen to users after they log in. The welcome screen needs to show the user's first name and the user's previous login date. According to company policy, developers who work on the web application cannot store any personally identifiable information in RDS instances.

Which solution should the developer implement to meet these requirements?

- A. After successful login, submit a Cognito request for user token
- B. When redirecting to the welcome screen, provide the identity token in the Authorization header of the reques
- C. Extract the user name from the given_name claim and the user's universally unique identifier (UUID) from the sub claim inside the identity toke
- D. Use the UUID as the key to store and retrieve the previous login information from the database.
- E. After successful login, submit a Cognito request for user token
- F. When redirecting to the welcome screen, provide the access token in the Authorization header of the reques
- G. Extract the user name from the given_name claim and the user's universally unique identifier (UUID) from the sub claim inside the access toke
- H. Use the UUID as the key to store and retrieve the previous login information from the database.
- I. After successful login, submit a Cognito request for user token
- J. When redirecting to the welcome screen, provide the identity token in the Authorization header of the reques
- K. Extract the user name from the given_name claim and the user's universally unique identifier (UUID) from the iss claim inside the identity toke
- L. Use the UUID as the key to store and retrieve the previous login information from the database.
- M. After successful login, submit a Cognito request for user token
- N. When redirecting to the welcome screen, provide the access token in the Authorization header of the reques
- O. Extract the user name from the given name claim and the user's universally unique identifier (UUID) from the iss claim inside the access toke
- P. Use the UUID as the key to store and retrieve the previous login information from the database.

Answer: C

NEW QUESTION 353

- (Exam Topic 4)

A media company wants to test its web application more frequently. The company deploys the application by using a separate AWS CloudFormation stack for each environment. The same CloudFormation template is deployed to each stack as the application progresses through the development lifecycle. A developer needs to build an automated alert for the quality assurance (QA) team. The developer wants the alert to occur for new deployments in the final pre-production environment.

Which solution will meet these requirements?

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic
- B. Add a subscription to notify the QA team
- C. Update the CloudFormation stack options to point to the SNS topic in the pre-production environment
- D. Most Voted
- E. Create an AWS Lambda function that notifies the QA team
- F. Create an Amazon EventBridge rule to invoke the Lambda function on the default event bus
- G. Filter the events on the CloudFormation service and the CloudFormation stack Amazon Resource Name (ARN).
- H. Create an Amazon CloudWatch alarm that monitors the metrics from CloudFormation
- I. Filter the metrics on the stack name and the stack status
- J. Configure the alarm to notify the QA team.
- K. Create an AWS Lambda function that notifies the QA team
- L. Configure the event source mapping to receive events from CloudFormation
- M. Specify the filtering values to limit invocations to the desired CloudFormation stack.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-rollback-email/>

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-rollback-email/> <https://www.trendmicro.com/cloudoneconformity/knowledge-base/aws/CloudFormation/cloudformation-stack-n>

NEW QUESTION 354

- (Exam Topic 4)

An open-source map application gathers data from several geolocation APIs. The application's source code repository is public and can be used by anyone, but the geolocation APIs must not be directly accessible.

A developer must implement a solution to prevent the credentials that are used to access the APIs from becoming public. The solution also must ensure that the application still functions properly.

Which solution will meet these requirements MOST cost-effectively?

- A. Store the credentials in AWS Secrets Manager
- B. Retrieve the credentials by using the GetSecretValue API operation.
- C. Store the credentials in AWS Key Management Service (AWS KMS). Retrieve the credentials by using the GetPublicKey API operation.
- D. Store the credentials in AWS Security Token Service (AWS STS). Retrieve the credentials by using the GetCallerIdentity API operation.
- E. Store the credentials in AWS Systems Manager Parameter Store
- F. Retrieve the credentials by using the GetParameter API operation.

Answer: D

Explanation:

Secrets Manager: It is paid. The storage cost is \$0.40 per secret per month and API interactions cost is \$0.05 per 10,000 API calls. Parameter Store: For Standard parameters, No additional charge for storage and standard throughput. For higher throughput, API interactions cost is \$0.05 per 10,000 API calls. For Advanced parameters, storage cost is \$0.05 per advanced parameter per month and API interactions cost is \$0.05 per 10,000 API calls.

<https://aws.amazon.com/systems-manager/pricing/>

NEW QUESTION 356

- (Exam Topic 4)

A company has a serverless application that uses AWS Lambda functions and AWS Systems Manager parameters to store configuration data. The company moves the Lambda functions inside the VPC and into private subnets. The Lambda functions are now producing errors in their attempts to access Systems Manager parameters.

Which solution will allow the Lambda functions to access Systems Manager parameters inside the VPC?

- A. Configure security groups to allow access to Systems Manager.
- B. Create an interface VPC endpoint for Systems Manager.
- C. Use an internet gateway from inside the VPC.
- D. Create a gateway VPC endpoint for Systems Manager.

Answer: B

NEW QUESTION 359

- (Exam Topic 4)

A distributed application includes an AWS Lambda function that runs successfully in the DEV environment with 128 MB of memory assigned. The same function is failing in the TEST environment. The developer is monitoring the application using AWS X-Ray, but the Lambda function cannot be seen on the X-Ray service graph. The Lambda execution role has AWS X-Ray permissions

What is the MOST LIKELY cause for AWS X-Ray not showing any data for the Lambda function?

- A. The AWS SDK needs to be included in the AWS Lambda deployment package.
- B. VPC Flow Logs are not enabled for the application VPC.
- C. Active tracing needs to be enabled for the Lambda function
- D. The memory needs to be increased to 2 GB for the TEST environments.

Answer: C

Explanation:

<https://stackoverflow.com/questions/43728674/enabling-x-ray-support-in-aws-lambda> You need to check the "Enable Active Tracing" checkbox in the Lambda console.

NEW QUESTION 361

- (Exam Topic 4)

A developer is migrating to Amazon Cognito from a custom user management solution that stores user information in a database. The developer has created a..... Amazon Cognito user pool. The developer needs to migrate the existing user information to the user pool without forcing users to change their passwords. Which solution will meet these requirements?

- A. Import users from a .csv file.
- B. Add an OpenID Connect (OIDC) identity provider to the user pool.
- C. Import users from a .json file.
- D. Import users with a user migration AWS Lambda trigger.

Answer: B

NEW QUESTION 364

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

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

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Developer-Associate Practice Test Here](#)