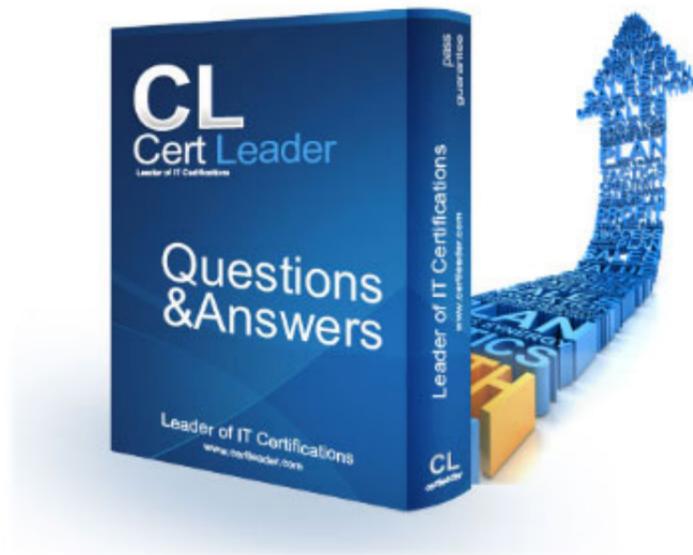


AWS-Certified-Developer-Associate Dumps

Amazon AWS Certified Developer - Associate

<https://www.certleader.com/AWS-Certified-Developer-Associate-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

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

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 2

- (Exam Topic 1)

An application stops working with the following error: The specified bucket does not exist. Where is the BEST place to start the root cause analysis?

- A. Check the Elastic Load Balancer logs for DeleteBucket requests.
- B. Check the application logs in Amazon CloudWatch Logs for Amazon S3 DeleteBucket errors.
- C. Check AWS X-Ray for Amazon S3 DeleteBucket alarms.
- D. Check AWS CloudTrail for a DeleteBucket event.

Answer: D

NEW QUESTION 3

- (Exam Topic 1)

A company has three different environments: Development, QA, and Production. The company wants to deploy its code first in the Development environment, then QA, and then Production.

Which AWS service can be used to meet this requirement?

- A. Use AWS CodeCommit to create multiple repositories to deploy the application.
- B. Use AWS CodeBuild to create, configure, and deploy multiple build application projects.
- C. Use AWS Data Pipeline to create multiple data pipeline provisions to deploy the application.
- D. Use AWS CodeDeploy to create multiple deployment groups.

Answer: D

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-groups.html>

"You can associate more than one deployment group with an application in CodeDeploy. This makes it possible to deploy an application revision to different sets of instances at different times. For example, you might use one deployment group to deploy an application revision to a set of instances tagged Test where you ensure the quality of the code. Next, you deploy the same application revision to a deployment group with instances tagged Staging for additional verification. Finally, when you are ready to release the latest application to customers, you deploy to a deployment group that includes instances tagged Production. "

NEW QUESTION 4

- (Exam Topic 1)

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

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

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

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

Answer: CE

NEW QUESTION 5

- (Exam Topic 1)

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

Which approach should the Developer use to handle this change?

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

Answer: C

NEW QUESTION 6

- (Exam Topic 1)

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

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

Answer: ACE

NEW QUESTION 7

- (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 8

- (Exam Topic 1)

What is the maximum number of S3 Buckets available per AWS account?

- A. 100 per region
- B. there is no limit
- C. 100 per account
- D. 500 per account
- E. 100 per IAM user

Answer: C

Explanation:

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

NEW QUESTION 9

- (Exam Topic 1)

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

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

Answer: E

Explanation:

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

NEW QUESTION 10

- (Exam Topic 1)

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

Which source control service should be used?

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

Answer: C

NEW QUESTION 10

- (Exam Topic 1)

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

How can the Developer resolve this issue?

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

Answer: C

Explanation:

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

NEW QUESTION 12

- (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 14

- (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 18

- (Exam Topic 1)

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

How would the Developer optimize this scan?

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

Answer: A

Explanation:

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

NEW QUESTION 22

- (Exam Topic 1)

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

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

Answer: A

NEW QUESTION 23

- (Exam Topic 1)

Company D is running their corporate website on Amazon S3 accessed from <http://www.companyd.com>. Their marketing team has published new web fonts to a separate S3 bucket accessed by the S3 endpoint

<https://s3-us-west-1.amazonaws.com/cdfonts>. While testing the new web fonts, Company D recognized the web fonts are being blocked by the browser.

What should Company D do to prevent the web fonts from being blocked by the browser?

- A. Enable versioning on the cdfonts bucket for each web font
- B. Create a policy on the cdfonts bucket to enable access to everyone
- C. Add the Content-MD5 header to the request for webfonts in the cdfonts bucket from the website
- D. Configure the cdfonts bucket to allow cross-origin requests by creating a CORS configuration

Answer: D

Explanation:

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

NEW QUESTION 27

- (Exam Topic 1)

An application running on Amazon EC2 instances must access objects within an Amazon S3 bucket that are encrypted using server-side encryption using AWS KMS encryption keys (SSE-KMS). The application must have access to the customer master key (CMK) to decrypt the objects.

Which combination of steps will grant the application access? (Select TWO.)

- A. Write an S3 bucket policy that grants the bucket access to the key.
- B. Grant access to the key in the IAM EC2 role attached to the application's EC2 instances.
- C. Write a key policy that enables IAM policies to grant access to the key.
- D. Grant access to the key in the S3 bucket's ACL
- E. Create a Systems Manager parameter that exposes the KMS key to the EC2 instances.

Answer: BC

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/decrypt-kms-encrypted-objects-s3/> IAM role needs access to the keys to decrypt the object and key policies must allow role access to the key. Key policies are the primary way to control access to customer master keys (CMKs) in AWS KMS. You need the permission to decrypt the AWS KMS key. When a user sends a GET request, Amazon S3 checks if the AWS Identity and Access Management (IAM) user or role that sent the request is authorized to decrypt the key associated with the object. If the IAM user or role belongs to the same AWS account as the key, then the permission to decrypt must be granted on the AWS KMS key's policy.

NEW QUESTION 32

- (Exam Topic 1)

EC2 instances are launched from Amazon Machine images (AMIs). A given public AMI can:

- A. be used to launch EC2 Instances in any AWS region.
- B. only be used to launch EC2 instances in the same country as the AMI is stored.
- C. only be used to launch EC2 instances in the same AWS region as the AMI is stored.
- D. only be used to launch EC2 instances in the same AWS availability zone as the AMI is stored

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/usingsharedamis-finding.html>

NEW QUESTION 33

- (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 35

- (Exam Topic 1)

Which of the following items are required to allow an application deployed on an EC2 instance to write data to a DynamoDB table?

Assume that no security Keys are allowed to be stored on the EC2 instance. Choose 2 answers

- A. Create an IAM User that allows write access to the DynamoDB table.
- B. Add an IAM Role to a running EC2 instance.
- C. Add an IAM User to a running EC2 Instance.
- D. Launch an EC2 Instance with the IAM Role included in the launch configuration.
- E. Create an IAM Role that allows write access to the DynamoDB table.
- F. Launch an EC2 Instance with the IAM User included in the launch configuration.

Answer: BE

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html#attach-iam-role>

NEW QUESTION 40

- (Exam Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 43

- (Exam Topic 1)

A Development team has pushed out 10 applications running on several Amazon EC2 instances. The Operations team is asking for a graphical representation of one key performance metric for each application. These metrics should be available on one screen for easy monitoring. Which steps should the Developer take to accomplish this using Amazon CloudWatch?

- A. Create a custom namespace with a unique metric name for each application.
- B. Create a custom dimension with a unique metric name for each application.
- C. Create a custom event with a unique metric name for each application.
- D. Create a custom alarm with a unique metric name for each application.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudwatch-custom-metrics/>

NEW QUESTION 48

- (Exam Topic 1)

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

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

Answer: B

NEW QUESTION 50

- (Exam Topic 1)

A company recently migrated its web, application and NoSQL database tiers to AWS. The company is using Auto Scaling to scale the web and application tiers. More than 95 percent of the Amazon DynamoDB requests are repeated read-requests. How can the DynamoDB NoSQL tier be scaled up to cache these repeated requests?

- A. Amazon EMR
- B. Amazon DynamoDB Accelerator
- C. Amazon SQS
- D. Amazon CloudFront

Answer: B

Explanation:

Reference: <https://aws.amazon.com/dynamodb/dax/>

NEW QUESTION 52

- (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 54

- (Exam Topic 1)

What is one key difference between an Amazon EBS-backed and an instance-store backed instance?

- A. Virtual Private Cloud requires EBS backed instances
- B. Amazon EBS-backed instances can be stopped and restarted
- C. Auto scaling requires using Amazon EBS-backed instances.
- D. Instance-store backed instances can be stopped and restarted.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/instance-store-vs-ebs/>

NEW QUESTION 58

- (Exam Topic 1)

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

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

Answer: ACD

Explanation:

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

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

NEW QUESTION 62

- (Exam Topic 1)

Company B provides an online image recognition service and utilizes SQS to decouple system components for scalability. The SQS consumers poll the imaging queue as often as possible to keep end-to-end throughput as high as possible. However, Company B is realizing that polling in tight loops is burning CPU cycles and increasing costs with empty responses.

How can Company B reduce the number of empty responses?

- A. Set the imaging queue visibility Timeout attribute to 20 seconds
- B. Set the Imaging queue ReceiveMessageWaitTimeSeconds attribute to 20 seconds
- C. Set the imaging queue MessageRetentionPeriod attribute to 20 seconds
- D. Set the DelaySeconds parameter of a message to 20 seconds

Answer: B

Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-short-and-long-polling>.

NEW QUESTION 63

- (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 66

- (Exam Topic 1)

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

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

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

Answer: C

NEW QUESTION 68

- (Exam Topic 1)

A company needs to encrypt data at rest, but it wants to leverage an AWS managed service using its own master key.

Which of the following AWS service can be used to meet these requirements?

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

Answer: B

NEW QUESTION 71

- (Exam Topic 1)

How should custom libraries be utilized in AWS Lambda?

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

Answer: D

Explanation:

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

NEW QUESTION 74

- (Exam Topic 1)

What does an Amazon SQS delay queue accomplish?

- A. Messages are hidden for a configurable amount of time when they are first added to the queue.
- B. Messages are hidden for a configurable amount of time after they are consumed from the queue.
- C. The consumer can poll the queue for a configurable amount of time before retrieving a message.
- D. Message cannot be deleted for a configurable amount of time after they are consumed from the queue.

Answer: A

Explanation:

Reference:

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

NEW QUESTION 78

- (Exam Topic 1)

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

Explanation:

[https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task_IAM_role.html#:~:targetText=Amazon%](https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task_IAM_role.html#:~:targetText=Amazon%20IAM%20role)

NEW QUESTION 81

- (Exam Topic 1)

A Development team currently supports an application that uses an in-memory store to save accumulated game results. Individual results are stored in a database. As part of migrating to AWS, the team needs to use automatic scaling. The team knows this will yield inconsistent results.

Where should the team store these accumulated game results to BEST allow for consistent results without impacting performance?

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

D. Amazon Kinesis

Answer: C

NEW QUESTION 83

- (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.
- 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 87

- (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 89

- (Exam Topic 1)

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

How can the Developer achieve these requirements?

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

Answer: D

NEW QUESTION 94

- (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 96

- (Exam Topic 1)

A company has a website that is developed in PHP and WordPress and is launched using AWS Elastic Beanstalk. There is a new version of the website that needs to be deployed in the Elastic Beanstalk environment. The company cannot tolerate having the website offline if an update fails. Deployments must have minimal impact and rollback as soon as possible. What deployment method should be used?

- A. All at once
- B. Rolling
- C. Snapshots
- D. Immutable

Answer: D

Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environmentmgmt-updates-immutable.html>

NEW QUESTION 101

- (Exam Topic 1)

An application stores images in an S3 bucket. Amazon S3 event notifications are used to trigger a Lambda function that resizes the images. Processing each image takes less than a second.

How will AWS Lambda handle the additional traffic?

- A. Lambda will scale out to execute the requests concurrently.
- B. Lambda will handle the requests sequentially in the order received.
- C. Lambda will process multiple images in a single execution.
- D. Lambda will add more compute to each execution to reduce processing time.

Answer: A

NEW QUESTION 102

- (Exam Topic 1)

A company is using AWS CodeBuild to compile a website from source code stored in AWS CodeCommit. A recent change to the source code has resulted in the CodeBuild project being unable to successfully compile the website.

How should the Developer identify the cause of the failures?

- A. Modify the buildspec.yml file to include steps to send the output of build commands to Amazon CloudWatch.
- B. Use a custom Docker image that includes the AWS X-Ray agent in the AWS CodeBuild project configuration.
- C. Check the build logs of the failed phase in the last build attempt in the AWS CodeBuild project build history.
- D. Manually re-run the build process on a local machine so that the output can be visualized.

Answer: A

NEW QUESTION 104

- (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 reference the associated IAM role.

Answer: D

Explanation:

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

NEW QUESTION 105

- (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 109

- (Exam Topic 1)

Which of the following programming languages have an officially supported AWS SDK? Choose 2 answers

- A. Perl
- B. PHP
- C. Pascal
- D. Java
- E. SQL

Answer: BD

NEW QUESTION 112

- (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 117

- (Exam Topic 1)

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

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

Answer: C

NEW QUESTION 121

- (Exam Topic 1)

The Lambda function below is being called through an API using Amazon API Gateway. The average execution time for the Lambda function is about 1 second. The pseudocode for the Lambda function is as shown in the exhibit.

```
include "3rd party encryption module"
include "match module"
lambda_handler(event, context)
    rds_host = "rds-instance-endpoint"
    name = db_username
    password = db_password
    db_name = db_name
    # Connect to the RDS Database
    Conn = RDSConnection(rds_host, user=name, passwd=password,
    db=db_name, connect_timeout=5)
    #Perform some Processing reading data from the RDS database
    #Code Block
    #Code Block
    #Code Block
```

What two actions can be taken to improve the performance of this Lambda function without increasing the cost of the solution? (Select two.)

- A. Package only the modules the Lambda function requires
- B. Use Amazon DynamoDB instead of Amazon RDS
- C. Move the initialization of the variable Amazon RDS connection outside of the handler function
- D. Implement custom database connection pooling with the Lambda function
- E. Implement local caching of Amazon RDS data so Lambda can re-use the cache

Answer: AC

NEW QUESTION 123

- (Exam Topic 1)

A Developer will be using the AWS CLI on a local development server to manage AWS services.

What can be done to ensure that the CLI uses the Developer's IAM permissions when making commands?

- A. Specify the Developer's IAM access key ID and secret access key as parameters for each CLI command.
- B. Run the aws configure CLI command, and provide the Developer's IAM access key ID and secret access key.
- C. Specify the Developer's IAM user name and password as parameters for each CLI command.
- D. Use the Developer's IAM role when making the CLI command.

Answer: B

Explanation:

<https://medium.com/faun/setting-up-a-production-environment-using-our-local-development-server-and-aws-f5e>

NEW QUESTION 125

- (Exam Topic 1)

A company is building a stock trading application that requires sub-millisecond latency in processing trading requests. Amazon DynamoDB is used to store all the trading data that is used to process each request. After load testing the application, the development team found that due to data retrieval times, the latency requirement is not satisfied. Because of sudden high spikes in the number of requests, DynamoDB read capacity has to be significantly over-provisioned to avoid throttling.

What steps should be taken to meet latency requirements and reduce the cost of running the application?

- A. Add Global Secondary Indexes for trading data.
- B. Store trading data in Amazon S3 and use Transfer Acceleration.
- C. Add retries with exponential back-off for DynamoDB queries
- D. Use DynamoDB Accelerator to cache trading data.

Answer: D

Explanation:

Refer AWS documentation - DynamoDB Accelerator

Amazon DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement – from milliseconds to microseconds – even at millions of requests per second. DAX does all the heavy lifting required to add in-memory acceleration to your DynamoDB tables, without requiring developers to manage cache invalidation, data population, or cluster management. Now you can focus on building great applications for your customers without worrying about performance at scale.

NEW QUESTION 129

- (Exam Topic 1)

A Developer wants access to make the log data of an application running on an EC2 instance available to systems administrators.

Which of the following enables monitoring of this metric in Amazon CloudWatch?

- A. Retrieve the log data from CloudWatch using the GetMetricData API call
- B. Retrieve the log data from AWS CloudTrail using the LookupEvents API call.
- C. Launch a new EC2 instance, configure Amazon CloudWatch Events, and then install the application.
- D. Install the Amazon CloudWatch Logs agent on the EC2 instance that the application is running on.

Answer: D

NEW QUESTION 132

- (Exam Topic 1)

A company has an application that logs all information to Amazon S3. Whenever there is a new log file, an AWS Lambda function is invoked to process the log files. The code works, gathering all of the necessary information. However, when checking the Lambda function logs, duplicate entries with the same request ID are found.

What is causing the duplicate entries?

- A. The S3 bucket name was specified incorrectly.
- B. The Lambda function failed, and the Lambda service retried the invocation with a delay.
- C. There was an S3 outage, which caused duplicate entries of the sale log file.
- D. The application stopped intermittently and then resumed.

Answer: B

Explanation:

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

NEW QUESTION 136

- (Exam Topic 1)

An organization must store thousands of sensitive audio and video files in an Amazon S3 bucket. Organizational security policies require that all data written to this bucket be encrypted.

How can compliance with this policy be ensured?

- A. Use AWS Lambda to send notifications to the security team if unencrypted objects are put in the bucket.
- B. Configure an Amazon S3 bucket policy to prevent the upload of objects that do not contain the x-amz-server-side-encryption header.
- C. Create an Amazon CloudWatch event rule to verify that all objects stored in the Amazon S3 bucket are encrypted.
- D. Configure an Amazon S3 bucket policy to prevent the upload of objects that contain the x-amz-server-side-encryption header.

Answer: B

NEW QUESTION 140

- (Exam Topic 1)

Which of the following services are key/value stores? Choose 3 answers

- A. Amazon ElastiCache
- B. Simple Notification Service
- C. DynamoDB
- D. Simple Workflow Service
- E. Simple Storage Service

Answer: ACE

NEW QUESTION 145

- (Exam Topic 1)

An existing serverless application processes uploaded image files. The process currently uses a single Lambda function that takes an image file, performs the

processing, and stores the file in Amazon S3. Users of the application now require thumbnail generation of the images. Users want to avoid any impact to the time it takes to perform the image uploads.

How can thumbnail generation be added to the application, meeting user requirements while minimizing changes to existing code?

- A. Change the existing Lambda function handling the uploads to create thumbnails at the time of upload. Have the function store both the image and thumbnail in Amazon S3.
- B. Create a second Lambda function that handles thumbnail generation and storage.
- C. Change the existing Lambda function to invoke it asynchronously.
- D. Create an S3 event notification with a Lambda function destination.
- E. Create a new Lambda function to generate and store thumbnails.
- F. Create an S3 event notification to an SQS Queue.
- G. Create a scheduled Lambda function that processes the queue, and generates and stores thumbnails.

Answer: C

Explanation:

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

NEW QUESTION 146

- (Exam Topic 1)

A company developed a set of APIs that are being served through the Amazon API Gateway. The API calls need to be authenticated based on OpenID identity providers such as Amazon or Facebook. The APIs should allow access based on a custom authorization model.

Which is the simplest and MOST secure design to use to build an authentication and authorization model for the APIs?

- A. Use Amazon Cognito user pools and a custom authorizer to authenticate and authorize users based on JSON Web Tokens.
- B. Build a OpenID token broker with Amazon and Facebook.
- C. Users will authenticate with these identity providers and pass the JSON Web Token to the API to authenticate each API call.
- D. Store user credentials in Amazon DynamoDB and have the application retrieve temporary credentials from AWS STS.
- E. Make API calls by passing user credentials to the APIs for authentication and authorization.
- F. Use Amazon RDS to store user credentials and pass them to the APIs for authentications and authorization.

Answer: A

NEW QUESTION 150

- (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 153

- (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 158

- (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 162

- (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 163

- (Exam Topic 1)

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

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

Answer: AE

Explanation:

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

"There are no limits to the number of prefixes in a bucket. You can increase your read or write performance by parallelizing reads. For example, if you create 10 prefixes in an Amazon S3 bucket to parallelize reads, you could scale your read performance to 55,000 read requests per second." The assumption that prefixes don't matter is incorrect, as described by "Amazon S3 performance guidelines recommended randomizing prefix naming with **hashed characters** to optimize performance for frequent data retrievals. You no longer have to randomize prefix naming for performance, and can use sequential date-based naming for your prefixes"

NEW QUESTION 167

- (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 168

- (Exam Topic 1)

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

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

- A. Use AWS KMS to store and retrieve credentials.
- B. Use EC2 instance profiles.

- C. Use AWS root user to make requests to the application.
- D. Store and retrieve credentials from AWS CodeCommit.

Answer: B

Explanation:

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

NEW QUESTION 172

- (Exam Topic 1)

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

What is the recommended solution?

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

Answer: D

NEW QUESTION 177

- (Exam Topic 1)

A Developer is creating a Lambda function and will be using external libraries that are not included in the standard Lambda libraries.

What action would minimize the Lambda compute time consumed?

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

Answer: D

NEW QUESTION 182

- (Exam Topic 2)

A social media company is using Amazon Cognito in order to synchronize profiles across different mobile devices, to enable end users to have a seamless experience.

Which of the following configurations can be used to silently notify users whenever an update is available on all other devices?

- A. Modify the user pool to include all the devices which keep them in sync.
- B. Use the SyncCallback interface to receive notifications on the application.
- C. Use an Amazon Cognito stream to analyze the data and push the notifications.
- D. Use the push synchronization feature with the appropriate IAM role.

Answer: D

Explanation:

<https://docs.aws.amazon.com/cognito/latest/developerguide/push-sync.html>

NEW QUESTION 183

- (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 187

- (Exam Topic 2)

A developer has written a serverless application and wants to deploy it to AWS Lambda to leverage the function's multi-threaded execution to improve performance. Which action should the developer take to achieve these requirements?

- A. increase the Lambda function execution timeout
- B. Use unreserved account concurrency.
- C. Increase the memory allocation of the Lambda function
- D. Set the reserved concurrency of the Lambda function to a higher number

Answer: C

NEW QUESTION 192

- (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 197

- (Exam Topic 2)

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

Which type of encryption meets these requirements?

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

Answer: B

NEW QUESTION 200

- (Exam Topic 2)

An application running on Amazon EC2 opens connections to an Amazon RDS SQL Server database. The developer does not want to store the user name and password for the database in the code. The developer would also like to automatically rotate the credentials.

What is the MOST secure way to store and access the database credentials?

- A. Create an IAM role that has permissions to access the database. Attach the role to the EC2 instance.
- B. Use AWS Secrets Manager to store the credential.
- C. Retrieve the credentials from Secrets Manager as needed.
- D. Store the credentials in an encrypted text file in an Amazon S3 bucket. Configure the EC2 instance's user data to download the credentials from Amazon S3 as the instance boots.
- E. Store the user name and password credentials directly in the source code.
- F. No further action is needed because the source code is stored in a private repository.

Answer: B

NEW QUESTION 203

- (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 208

- (Exam Topic 2)

A web application is designed to allow new users to create accounts using their email addresses. The application will store attributes for each user, and is expecting millions of users to sign up.

What should the Developer implement to achieve the design goals?

- A. Amazon Cognito user pools
- B. AWS Mobile Hub user data storage
- C. Amazon Cognito Sync
- D. AWS Mobile Hub cloud logic

Answer: A

NEW QUESTION 210

- (Exam Topic 2)

A developer is provided with an HTTPS clone URL for an AWS CodeCommit repository. What needs to be configured before cloning this repository?

- A. Use AWS KMS to set up public and private keys for use with CodeCommit.
- B. Set up the Git credential helper to use an AWS credential profile, and enable the helper to send the path to the repositories.
- C. Generate encryption keys using AWS CloudHSM, then export the key for use with AWS CodeCommit.
- D. Use AWS certificate manager to provision public and private SSL/TLS certificates.

Answer: B

Explanation:

AWS credential profile, and enabling the Git credential helper to send the path to repositories: Reference:
<https://docs.aws.amazon.com/codecommit/latest/userguide/setting-up-https-unixes.html>

NEW QUESTION 215

- (Exam Topic 2)

A Developer is working on an application that handles 10MB documents that contain highly-sensitive data. The application will use AWS KMS to perform client-side encryption.

What steps must be followed?

- A. Invoke the Encrypt API passing the plaintext data that must be encrypted, then reference the customer managed key ARN in the KeyId parameter
- B. Invoke the GenerateRandom API to get a data encryption key, then use the data encryption key to encrypt the data
- C. Invoke the GenerateDataKey API to retrieve the encrypted version of the data encryption key to encrypt the data
- D. Invoke the GenerateDataKey API to retrieve the plaintext version of the data encryption key to encrypt the data

Answer: D

Explanation:

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

GenerateDataKey API: Generates a unique data key. This operation returns a plaintext copy of the data key and a copy that is encrypted under a customer master key (CMK) that you specify. You can use the plaintext key to encrypt your data outside of KMS and store the encrypted data key with the encrypted data.

NEW QUESTION 217

- (Exam Topic 2)

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

What AWS service should be used to accomplish this?

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

Answer: A

Explanation:

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

References:

NEW QUESTION 221

- (Exam Topic 2)

A Developer has developed a web application and wants to deploy it quickly on a Tomcat server on AWS. The Developer wants to avoid having to manage the underlying infrastructure.

What is the easiest way to deploy the application, based on these requirements?

- A. AWS CloudFormation
- B. AWS Elastic Beanstalk
- C. Amazon S3
- D. AWS CodePipeline

Answer: B

NEW QUESTION 225

- (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 230

- (Exam Topic 2)

A company is developing a web application that allows its employees to upload a profile picture to a private Amazon S3 bucket. There is no size limit for the profile pictures, which should be displayed every time an employee logs in. For security reasons, the pictures cannot be publicly accessible.

What is a viable long-term solution for this scenario?

- A. Generate a presigned URL when a picture is uploaded. Save the URL in an Amazon DynamoDB table. Return the URL to the browser when the employee logs in.
- B. Save the picture's S3 key in an Amazon DynamoDB table. Create an Amazon S3 VPC endpoint to allow the employees to download pictures once they log in.
- C. Encode a picture using base64. Save the base64 string in an Amazon DynamoDB table. Allow the browser to retrieve the string and convert it to a picture.

- D. Save the picture's S3 key in an Amazon DynamoDB tabl
- E. Use a function to generate a presigned URL every time an employee logs i
- F. Return the URL to the browser.

Answer: B

Explanation:

Reference:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-private-connection-noauthentication/>

NEW QUESTION 233

- (Exam Topic 2)

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

Which caching strategy will address this situation MOST efficiently?

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

Answer: B

Explanation:

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

NEW QUESTION 238

- (Exam Topic 2)

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

Why is the table being throttled?

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

Answer: A

Explanation:

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

NEW QUESTION 240

- (Exam Topic 2)

An application deployed on AWS Elastic Beanstalk experiences increased error rates during deployments of new application versions, resulting in service degradation for users. The Development team believes that this is because of the reduction in capacity during the deployment steps. The team would like to change the deployment policy configuration of the environment to an option that maintains full capacity during deployment while using the existing instances. Which deployment policy will meet these requirements while using the existing instances?

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

Answer: D

Explanation:

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

NEW QUESTION 243

- (Exam Topic 2)

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

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

- H. Alter the application code to inspect a custom header
- I. Alter the client code to pass the IP address in the custom header.

Answer: C

NEW QUESTION 245

- (Exam Topic 2)

A developer is testing an application that invokes an AWS Lambda function asynchronously. During the testing phase, the Lambda function fails to process after two retries.

How can the developer troubleshoot the failure?

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

Answer: A

NEW QUESTION 249

- (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 253

- (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 257

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 260

- (Exam Topic 2)

A developer has a legacy application that is hosted on-premises. Other applications hosted on AWS depend on the on-premises application for proper functioning. In case of any application errors, the developer wants to be able to use Amazon CloudWatch to monitor and troubleshoot all applications from one place.

How can the developer accomplish this?

- A. Install an AWS SDK on the on-premises server to automatically send logs to CloudWatch.

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 261

- (Exam Topic 2)

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

Which solution BEST meets these requirements?

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

Answer: C

NEW QUESTION 266

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 271

- (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 272

- (Exam Topic 2)

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

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

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

Answer: A

NEW QUESTION 276

- (Exam Topic 2)

A Developer needs to deploy an application running on AWS Fargate using Amazon ECS. The application has environment variables that must be passed to a container for the application to initialize

How should the environment variables be passed to the container?

- A. Define an array that includes the environment variables under the environment parameter within the service definition
- B. Define an array that includes the environment variables under the environment parameter within the task definition
- C. Define an array that includes the environment variables under the entrypoint parameter within the task definition
- D. Define in array that includes the environment variables under the EntryPoint parameter within the service definition

Answer: B

NEW QUESTION 277

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 280

- (Exam Topic 2)

A Developer wants to upload data to Amazon S3 and must encrypt the data in transit. Which of the following solutions will accomplish this task? (Choose two.)

- A. Set up hardware VPN tunnels to a VPC and access S3 through a VPC endpoint
- B. Set up Client-Side Encryption with an AWS KMS-Managed Customer Master Key
- C. Set up Server-Side Encryption with AWS KMS-Managed Keys
- D. Transfer the data over an SSL connection
- E. Set up Server-Side Encryption with S3-Managed Keys

Answer: BD

Explanation:

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

NEW QUESTION 283

- (Exam Topic 2)

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

What solution would give the BEST performance increase?

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

Answer: A

Explanation:

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

NEW QUESTION 288

- (Exam Topic 2)

A company is using AWS CodePipeline to deliver one of its applications. The delivery pipeline is triggered by changes to the master branch of an AWS CodeCommit repository and uses AWS CodeBuild to implement the test and build stages of the process and AWS CodeDeploy to deploy the application. The pipeline has been operating successfully for several months and there have been no modifications. Following a recent change to the application's source code, AWS CodeDeploy has not deployed the updates application as expected.

What are the possible causes? (Choose two.)

- A. The change was not made in the master branch of the AWS CodeCommit repository.
- B. One of the earlier stages in the pipeline failed and the pipeline has terminated.
- C. One of the Amazon EC2 instances in the company's AWS CodePipeline cluster is inactive.
- D. The AWS CodePipeline is incorrectly configured and is not executing AWS CodeDeploy.
- E. AWS CodePipeline does not have permissions to access AWS CodeCommit.

Answer: AB

NEW QUESTION 291

- (Exam Topic 2)

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

Which AWS service should be used to accomplish this?

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

Answer: D

NEW QUESTION 295

- (Exam Topic 2)

A company has implemented AWS CodePipeline to automate its release pipelines. The development team is writing an AWS Lambda function that will send notifications for state changes of each of the actions in the stages.

Which steps must be taken to associate the Lambda function with the event source?

- A. Create a trigger that invokes the Lambda function from the Lambda console by selecting CodePipeline as the event source
- B. Create an event trigger and specify the Lambda function from the CodePipeline console.
- C. Create an Amazon CloudWatch alarm that monitors status changes in CodePipeline and triggers the Lambda function
- D. Create an Amazon CloudWatch Events rule that uses CodePipeline as an event source.

Answer: B

NEW QUESTION 300

- (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 status code, status description, and headers in the Lambda function
- D. Enable the multi-value headers on the ALB

Answer: D

NEW QUESTION 304

- (Exam Topic 2)

A developer is writing a web application that must share secure documents with end users. The documents are stored in a private Amazon S3 bucket. The application must allow only authenticated users to download specific documents when requested, and only for a duration of 15 minutes.

How can the developer meet these requirements?

- A. Copy the documents to a separate S3 bucket that has a lifecycle policy for deletion after 15 minutes
- B. Create a presigned S3 URL using the AWS SDK with an expiration time of 15 minutes
- C. Use server-side encryption with AWS KMS managed keys (SSE-KMS) and download the documents using HTTPS
- D. Modify the S3 bucket policy to only allow specific users to download the documents. Revert the change after 15 minutes.

Answer: B

NEW QUESTION 307

- (Exam Topic 2)

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

What should the developer do to update the application to support data storage in Amazon RDS?

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

Answer: A

NEW QUESTION 311

- (Exam Topic 2)

A company has developed a new serverless application using AWS Lambda functions that will be deployed using the AWS Serverless Application Model (AWS SAM) CLI. Which step should the developer complete prior to deploying the application?

- A. Compress the application to a .zip file and upload it into AWS Lambda
- B. Test the new AWS Lambda function by first tracing it in AWS X-Ray.
- C. Bundle the serverless application using a SAM package
- D. Create the application environment using the `eb create my-env` command.

Answer: B

NEW QUESTION 314

- (Exam Topic 2)

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

Which solution will meet this requirement?

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

Answer: C

NEW QUESTION 317

- (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 318

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 323

- (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 326

- (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 331

- (Exam Topic 2)

A company is running a Docker application on Amazon ECS. The application must scale based on user load in the last 15 seconds. How should a Developer instrument the code so that the requirement can be met?

- A. Create a high-resolution custom Amazon CloudWatch metric for user activity data, then publish data every 30 seconds
- B. Create a high-resolution custom Amazon CloudWatch metric for user activity data, then publish data every 5 seconds

- C. Create a standard-resolution custom Amazon CloudWatch metric for user activity data, then publish data every 30 seconds
- D. Create a standard-resolution custom Amazon CloudWatch metric for user activity data, then publish data every 5 seconds

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html#high-resolution-m>

NEW QUESTION 333

- (Exam Topic 2)

A developer is writing an application in AWS Lambda To simplify testing and deployments, the developer needs the database connection string to be easily changed without modifying the Lambda code.

How can this requirement be met?

- A. Store the connection string as a secret in AWS Secrets Manager
- B. Store the connection string in an IAM user account.
- C. Store the connection string in AWS KMS
- D. Store the connection string as a Lambda layer.

Answer: A

NEW QUESTION 336

- (Exam Topic 2)

A developer is creating as AWS lambda function that generates a new file each time it runs. Each new file must be checked into an AWS CodeCommit repository hosted in the same AWS account.

How should the developer accomplish this?

- A. When the Lambda function starts, use the Git CLI to Clone the repositior
- B. Check the new file into the cloned repository and push the change.
- C. After the new file is created in Lambda, use cURL to invoke the CodeCommit AP
- D. Send the file to the repository.
- E. Use an AWS SDK to instantiate a CodeCommit Clie
- F. Invoke the put _ file method to add the file to the repository.
- G. Upload the new file to an Amazon S3 bucke
- H. Create an AWS step Function to accept S3 event
- I. In the step Function, add the new file to the repository.

Answer: D

NEW QUESTION 337

- (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 versio
- 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 cod
- D. Update the alias to direct 10% of users to the newly deployed versio
- 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 cod
- 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 340

- (Exam Topic 2)

A Developer is publishing critical log data to a log group in Amazon CloudWatch Lops, which was created 2 months ago. The Developer must encrypt the log data using an AWS KMS customer master key (CMK) so future data can be encrypted to comply with me company's security policy

How can the Developer meet this requirement?

- A. Use the Cloud Watch Logs console and enable the encrypt feature on the log group.
- B. Use the AWS CLI create-log-group command and specify the key Amazon Resource Name (ARN)
- C. Use the KMs console and associate the CMK with the log group
- D. Use the AWS CLI associate-Kms-key command and specify the key Amazon Resource Name (ARN)

Answer: C

NEW QUESTION 342

- (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 347

- (Exam Topic 2)

AWS CodeBuild builds code for an application, creates the Docker image, pushes the image to Amazon Elastic Container Registry (Amazon ECR), and tags the image with a unique identifier.

If the Developers already have AWS CLI configured on their workstations, how can the Docker images be pulled to the workstations?

- A. Run the following:docker pull REPOSITORY URI : TAG
- B. Run the output of the following:aws ecr get-loginand then run:docker pull REPOSITORY URI : TAG
- C. Run the following:aws ecr get-loginand then run:docker pull REPOSITORY URI : TAG
- D. Run the output of the following:aws ecr get-download-url-for-layerand then run:docker pull REPOSITORY URI : TAG

Answer: B

Explanation:

<https://docs.aws.amazon.com/cli/latest/reference/ecr/get-login.html>

NEW QUESTION 348

- (Exam Topic 2)

A gaming company is developing a mobile game application for iOS® and Android® platforms. This mobile game securely stores user data locally on the device. The company wants to allow users to use multiple device for the game, which requires user data synchronization across device.

Which service should be used to synchronize user data across devices without the need to create a backend application?

- A. AWS Lambda
- B. Amazon S3
- C. Amazon DynamoDB
- D. Amazon Cognito

Answer: D

NEW QUESTION 353

- (Exam Topic 2)

A company has a two-tier application running on an Amazon EC2 server that handles all of its AWS based e-commerce activity. During peak times, the backend servers that process orders are overloaded with requests. This results in some orders failing to process. A developer needs to create a solution that will re-factor the application.

Which steps will allow for more flexibility during peak times, while still remaining cost-effective? (Select TWO.)

- A. Increase the backend T2 EC2 instance size to xl to handle the largest possible load throughout the year
- B. implement an Amazon SQS queue to decouple the front-end and backend servers
- C. Use an Amazon SNS queue to decouple the front-end and backend servers.
- D. Migrate the backend servers to on-premises and pull from an Amazon SNS queue
- E. Modify the backend servers to pull from an Amazon SQS queue.

Answer: CD

NEW QUESTION 354

- (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 357

- (Exam Topic 3)

A Lambda function processes data before sending it to a downstream service. Each piece of data is approximately 1 MB in size. After a security audit, the function is now required to encrypt the data before sending it downstream. Which API call is required to perform the encryption?

- A. Pass the data to the KMS ReEncrypt API for encryption
- B. Use the KMS GenerateDataKey API to get an encryption key
- C. Use the KMS GenerateDataKeyWithoutPlainText API to get an encryption key
- D. Pass the data to KMS as part of the Encrypt API for encryption

Answer: D

NEW QUESTION 361

- (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 362

- (Exam Topic 3)

A company hosts a client-side web application for one of its subsidiaries on Amazon S3. The web application can be accessed through Amazon CloudFront from <https://www.example.com>. After a successful rollout, the company wants to host three more client-side web applications for its remaining subsidiaries on three separate S3 buckets.

To achieve this goal, a developer moves all the common JavaScript files and web fonts to a central S3 bucket that serves the web applications. However, during testing, the developer notices that the browser blocks the JavaScript files and web fonts.

What should the developer do to prevent the browser from blocking the JavaScript files and web fonts?

- A. Create four access points that allow access to the central S3 bucket.
- B. Assign an access point to each web application bucket.
- C. Create a bucket policy that allows access to the central S3 bucket.
- D. Attach the bucket policy to the central S3 bucket.
- E. Create a cross-origin resource sharing (CORS) configuration that allows access to the central S3 bucket. Add the CORS configuration to the central S3 bucket.
- F. Create a Content-MD5 header that provides a message integrity check for the central S3 bucket.
- G. Insert the Content-MD5 header for each web application request.

Answer: C

NEW QUESTION 367

- (Exam Topic 3)

A developer is changing the configuration for a CPU-intensive AWS Lambda function that runs once an hour. The function usually takes 45 seconds to run, but sometimes the run time is up to 1 minute. The timeout parameter is set to 3 minutes, and all other parameters are set to default.

The developer needs to optimize the run time of this function. Which solution will meet this requirement?

- A. Redeploy the function within the default VPC.
- B. Increase the function's memory.
- C. Redeploy the function by using Lambda layers.
- D. Increase the function's reserved concurrency.

Answer: B

NEW QUESTION 371

- (Exam Topic 3)

A three-tier application hosted on AWS uses Amazon RDS for MySQL as its database. A developer must ensure the database credentials are stored and accessed securely.

What is the MOST secure way for the developer to achieve this?

- A. Store the credentials in a configuration file and commit it to the GIT repository.
- B. Store the credentials in AWS Secrets Manager and enable automatic secret rotation.
- C. Store the credentials using Amazon RDS and enable automatic rotation.
- D. Store the credentials in code and handle credentials rotation within the application.

Answer: A

NEW QUESTION 373

- (Exam Topic 3)

A developer is working with a Docker application that needs to be quickly deployed using AWS without changing the infrastructure or configuring health checks. The application should be configured so that changes and updates can be made automatically without any downtime. Which solution will meet these requirements?

- A. Use AWS Elastic Beanstalk for application deployment and select an all-at-once update policy.
- B. Use AWS Elastic Beanstalk for application deployment and select a rolling deployment policy.
- C. Deploy the Docker container on an Amazon EC2 instance in an Auto Scaling group and configure a health check on the EC2 instance
- D. Deploy the Docker container using AWS Lambda and enable Amazon CloudWatch monitoring

Answer: A

NEW QUESTION 375

- (Exam Topic 3)

A company is adding items to an Amazon DynamoDB table from an AWS Lambda function that is written in Python. A developer needs to implement a solution that inserts records in the DynamoDB table and performs automatic retry when the insert fails. Which solution meets these requirements with MINIMUM code changes?

- A. Configure the Python code to run the AWS CLI through shell to call the PutItem operation
- B. Call the PutItem operation from Python by using the DynamoDB HTTP API
- C. Queue the items in AWS Glue, which will put them into the DynamoDB table
- D. Use the AWS software development kit (SDK) for Python (boto3) to call the PutItem operation

Answer: D

NEW QUESTION 377

- (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 380

- (Exam Topic 3)

A developer has written an application that uses Amazon API Gateway and AWS Lambda. The developer needs to configure the application so that the developer can visualize the application's components and identify performance bottlenecks. What should the developer do to meet these requirements?

- A. Enable AWS X-Ray tracing on the API Gateway stage
- B. Enable AWS X-Ray tracing on the API Gateway methods
- C. Enable Amazon CloudWatch Logs for API Gateway
- D. Enable Amazon CloudWatch Logs for Lambda

Answer: A

NEW QUESTION 381

- (Exam Topic 3)

A company is using Amazon API Gateway to manage its public-facing API. The CISO requires that the APIs be used by test account users only. What is the MOST secure way to restrict API access to users of this particular AWS account?

- A. Client-side SSL certificates for authentication
- B. API Gateway resource policies
- C. Cross-origin resource sharing (CORS)
- D. Usage plans

Answer: D

NEW QUESTION 383

- (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 385

- (Exam Topic 3)

A developer is trying to monitor an application's status by running a cron job that returns 1 if the service is up and 0 if the service is down. The developer created code that uses an AWS CLI put-metric-alarm command to publish the custom metrics to Amazon CloudWatch and create an alarm. However, the developer is unable to create an alarm as the custom metrics do not appear in the CloudWatch console. What is causing this issue?

- A. Sending custom metrics using the CLI is not supported
- B. The developer needs to use the put-metric-data command.
- C. The developer must use a unified CloudWatch agent to publish custom metrics
- D. The code is not running on an Amazon EC2 instance

Answer: B

NEW QUESTION 389

- (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 390

- (Exam Topic 3)

A developer at a company writes an AWS CloudFormation template. The template refers to subnets that were created by a separate AWS CloudFormation template that the company's network team wrote. When the developer attempts to launch the stack for the first time, the launch fails.

Which template coding mistakes could have caused this failure? (Select TWO.)

- A. The developer's template does not use the Ref intrinsic function to refer to the subnets
- B. The developer's template does not use the ImportValue intrinsic function to refer to the subnets
- C. The Mappings section of the developer's template does not refer to the subnets.
- D. The network team's template does not export the subnets in the Outputs section
- E. The network team's template does not export the subnets in the Mappings section

Answer: BD

NEW QUESTION 391

- (Exam Topic 3)

A company processes incoming documents from an Amazon S3 bucket. Users upload documents to an S3 bucket using a web user interface. Upon receiving files in S3, an AWS Lambda function is invoked to process the files, but the Lambda function times out intermittently.

If the Lambda function is configured with the default settings, what will happen to the S3 event when there is a timeout exception?

- A. Notification of a failed S3 event is sent as an email through Amazon SNS.
- B. The S3 event is sent to the default Dead Letter Queue.
- C. The S3 event is processed until it is successful.
- D. The S3 event is discarded after the event is retried twice.

Answer: D

NEW QUESTION 396

- (Exam Topic 3)

A developer is building a serverless application using AWS Lambda and must create a REST API using an HTTP GET method. What needs to be defined to meet this requirement? (Select TWO)

- A. A Lambda@Edge function
- B. An Amazon API Gateway with a Lambda function
- C. An exposed GET method in an Amazon API Gateway ID.
- D. An exposed GET method in the Lambda function
- E. An exposed GET method in Amazon Route 53

Answer: BE

NEW QUESTION 400

- (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 405

- (Exam Topic 3)

A company is using continuous integration/continuous deliver (CI/CD) system. A developer must automate the deployment of an application software package to Amazon EC2 instances and virtual servers that run on premises.

Which AWS services 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 406

- (Exam Topic 3)

A developer is building a WebSocket API using Amazon API Gateway. The payload sent to this API is JSON that includes an action key. This key can have three different values: create, update, and remove. The developer must integrate with different routes based on the value of the action key of the incoming JSON payload.

How can the developer accomplish this task with the LEAST amount of configuration?

- A. Deploy the WebSocket API to three stages for the respective routes: create, update, and remove
- B. Create a new route key and set the name as action
- C. Set the value of the route selection expression to action
- D. Set the value of the route selection expression to `$request.body.action`

Answer: D

NEW QUESTION 407

- (Exam Topic 3)

A video-hosting website has two types of members: those who pay a fee, and those who do not. Each video upload places a message in Amazon SQS. A fleet of Amazon EC2 instances polls Amazon SQS and processes each video.

The developer needs to ensure that the videos uploaded by the paying members are processed first. How can the developer meet this requirement?

- A. Create two SQS queues: one for paying members, and one for non-paying members. Poll the paying member queue first and then poll the non-paying member queue.
- B. Use SQS to set priorities on individual items within a single queue: give the paying members' videos the highest priority.
- C. Use SQS to set priorities on individual items within a single queue and use Amazon SNS to encode the videos.
- D. Create two Amazon SNS topics: one for paying members and one for non-paying members. Use SNS topic subscription priorities to differentiate between the two types of members.

Answer: B

NEW QUESTION 412

- (Exam Topic 3)

A developer has built an application using Amazon Cognito for authentication and authorization. After a user is successfully logged in to the application, the application creates a user record in an Amazon DynamoDB table.

What is the correct flow to authenticate the user and create a record in the DynamoDB table?

- A. Authenticate and get a token from an Amazon Cognito user pool.
- B. Use the token to access DynamoDB.
- C. Authenticate and get a token from an Amazon Cognito identity pool.
- D. Use the token to access DynamoDB.
- E. Authenticate and get a token from an Amazon Cognito user pool. Exchange the token for AWS credentials with an Amazon Cognito identity pool.
- F. Use the credential to access DynamoDB.
- G. Authenticate and get a token from an Amazon Cognito identity pool.
- H. Exchange the token for AWS credentials with an Amazon Cognito user pool.
- I. Use the credentials to access DynamoDB.

Answer: B

NEW QUESTION 413

- (Exam Topic 3)

A developer needs to deploy a new version to an AWS Elastic Beanstalk application. How can the developer accomplish this task?

- A. Upload and deploy the new application version in the Elastic Beanstalk console.
- B. Use the `eb init` CLI command to deploy a new version.
- C. Terminate the current Elastic Beanstalk environment and create a new one.
- D. Modify the `ebextensions` folder to add a source option to services.

Answer: A

NEW QUESTION 417

- (Exam Topic 3)

A developer has created a REST API using Amazon API Gateway. The developer wants to log who and how each caller accesses the API. The developer also wants to control how long the logs are kept. What should the developer do to meet these requirements?

- A. Enable API Gateway execution logging. Delete old logs using API Gateway retention settings.
- B. Enable API Gateway access logs. Use Amazon CloudWatch retention settings to delete old logs.
- C. Enable detailed Amazon CloudWatch metrics. Delete old logs with a recurring AWS Lambda function.
- D. Create and use API Gateway usage plan.
- E. Delete old logs with a recurring AWS Lambda function.

Answer: A

NEW QUESTION 421

- (Exam Topic 3)

A developer is working on a web application that runs on Amazon Elastic Container Service (Amazon ECS) and uses an Amazon DynamoDB table to store data. The application performs a large number of read requests against a small set of the table data. How can the developer improve the performance of these requests? (Select TWO)

- A. Create an Amazon ElastiCache cluster. Configure the application to cache data in the cluster.
- B. Create a DynamoDB Accelerator (DAX) cluster. Configure the application to use the DAX cluster for DynamoDB requests.
- C. Configure the application to make strongly consistent read requests against the DynamoDB table.
- D. Increase the read capacity of the DynamoDB table.
- E. Enable DynamoDB adaptive capacity.

Answer: AD

NEW QUESTION 426

- (Exam Topic 3)

A team deployed an AWS CloudFormation template to update a stack that already included an Amazon RDS DB instance. However, before the deployment of the update, the team changed the name of the DB instance on the template by mistake. The DeletionPolicy attribute for all resources was not changed from the default values.

What will be the result of this mistake?

- A. AWS CloudFormation will create a new database and delete the old one.
- B. AWS CloudFormation will create a new database and keep the old one.
- C. AWS CloudFormation will overwrite the existing database and rename it.
- D. AWS CloudFormation will leave the existing database and will not create a new one.

Answer: A

NEW QUESTION 428

- (Exam Topic 3)

A developer has created an AWS Lambda function that is written in Python. The Lambda function reads data from objects in Amazon S3 and writes data to an Amazon DynamoDB table.

The function is successfully invoked from an S3 event notification when an object is created. However, the function fails when it attempts to write to the DynamoDB table. What is the MOST likely cause of this issue?

- A. The Lambda function's concurrency limit has been exceeded.
- B. The DynamoDB table requires a global secondary index (GSI) to support writes.
- C. The Lambda function does not have IAM permissions to write to DynamoDB.
- D. The DynamoDB table is not running in the same Availability Zone as the Lambda function.

Answer: C

NEW QUESTION 431

- (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 436

- (Exam Topic 3)

A developer is planning to use an Amazon API Gateway and AWS Lambda to provide a REST API. The developer will have three distinct environments to manage: development, test, and production. How should the application be deployed while minimizing the number of resources to manage?

- A. Create a separate API Gateway and separate Lambda function for each environment in the same Region.
- B. Assign a Region for each environment and deploy API Gateway and Lambda to each Region.
- C. Create one API Gateway with multiple stages with one Lambda function with multiple aliases.
- D. Create one API Gateway and one Lambda function, and use a REST parameter to identify the environment.

Answer: C

NEW QUESTION 440

- (Exam Topic 3)

An application is using a custom library to make HTTP calls directly to AWS service endpoints. The application is experiencing transient errors that are causing processes to stop when each error is first encountered. A request has been made to make the application more resilient by adding error retries and exponential backoff.

How should a developer implement the changes with MINIMAL custom code?

- A. Add a Retry-After HTTP header to API requests.
- B. Use the AWS CLI to configure the retry settings in a named profile
- C. Change the custom library to retry on 5xx errors only
- D. Use an AWS SDK and set retry-specific configurations.

Answer: D

NEW QUESTION 442

- (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 445

- (Exam Topic 3)

A developer supports an application that accesses data in an Amazon DynamoDB table. One of the item attributes is expirationDate in the timestamp format. The application uses this attribute to find items, archive them, and remove them from the table based on the timestamp value.

The application will be decommissioned soon, and the developer must find another way to implement this functionality. The developer needs a solution that will require the least amount of code to write.

Which solution will meet these requirements?

- A. Enable TTL on the expirationDate attribute in the table.
- B. Create a DynamoDB stream.
- C. Create an AWS Lambda function to process the deleted items.
- D. Create a DynamoDB trigger for the Lambda function.
- E. Create two AWS Lambda functions: one to delete the items and one to process the items. Create a DynamoDB stream. Use the DeleteItem API operation to delete the items based on the expirationDate attribute. Use the GetRecords API operation to get the items from the DynamoDB stream and process them.
- F. Create two AWS Lambda functions: one to delete the items and one to process the items. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled rule to invoke the Lambda functions. Use the DeleteItem API operation to delete the items based on the expirationDate attribute. Use the GetRecords API operation to get the items from the DynamoDB table and process them.
- G. Enable TTL on the expirationDate attribute in the table. Specify an Amazon Simple Queue Service (Amazon SQS) dead-letter queue as the target to delete the items. Create an AWS Lambda function to process the items.

Answer: C

NEW QUESTION 449

- (Exam Topic 3)

A developer wants to modify the following AWS CloudFormation template to embed another CloudFormation stack:

```
{
  "AWSTemplateFormatVersion" : "2010-09-09",
  "Resources" : {
    "cfStack" : {
      "Properties" : {
        "TemplateURL" : "https://s3.amazonaws.com/cloudformation-templates/cf.template",
        "Parameters" : {
          "InstanceType" : "t3.small"
        }
      }
    }
  }
}
```

Which syntax should the developer add to the blank line of the CloudFormation template to meet this requirement?

- A. "Mapping" : "AWS::CloudFormation::Stack",
- B. "Type" : "AWS::CloudFormation::NestedStack",

- C. "Typ©" : "AWS::CloudFormation::Stack",
D. "Mapping" : "AWS::CloudFormation::NestedStack",

Answer: A

NEW QUESTION 450

- (Exam Topic 3)

A company is launching a new web application in the AWS Cloud. The company's development team is using AWS Elastic Beanstalk for deployment and maintenance. According to the company's change management process, the development team must evaluate changes for a specific time period before completing the rollout.

Which deployment policy meets this requirement?

- A. Immutable
- B. Rolling
- C. Rolling with additional batch
- D. Traffic splitting

Answer: A

NEW QUESTION 455

- (Exam Topic 3)

A developer is creating an application to process a large number of requests. Requests must be processed in order, and each request should be processed only once. How should Amazon SQS be deployed to achieve this?

- A. Configure First in First out (FIFO) delivery in a standard Amazon SQS queue to process requests.
- B. Use an SQS FIFO queue to process requests.
- C. Use the SetOrder attribute to ensure sequential request processing.
- D. Convert the standard queue to a FIFO queue by renaming the queue to use the fifo suffix.

Answer: B

NEW QUESTION 458

- (Exam Topic 3)

A company has a web application in an Amazon Elastic Container Service (Amazon ECS) cluster running hundreds of secure services in AWS Fargate containers. The services are in target groups routed by an Application Load Balancer (ALB). Application users log in to the website anonymously, but they must be authenticated using any OpenID Connect protocol-compatible identity provider (IdP) to access the secure services. Which authentication approach would meet these requirements with the LEAST amount of effort?

- A. Configure the services to use Amazon Cognito.
- B. Configure the ALB to use Amazon Cognito.
- C. Configure the services to use AWS Security Token Service (AWS STS) with the OpenID Connect IdP.
- D. Configure the Amazon ECS cluster to use AWS Security Token Service (AWS STS) with the OpenID Connect IdP.

Answer: A

NEW QUESTION 463

- (Exam Topic 3)

A developer is working on an ecommerce website. The developer wants to review server logs without logging in to each of the application servers individually. The website runs on multiple Amazon EC2 instances, is written in Python, and needs to be highly available.

How can the developer update the application to meet these requirements with MINIMUM changes?

- A. Rewrite the application to be cloud native and to run on AWS Lambda where the logs can be reviewed in Amazon CloudWatch.
- B. Set up centralized logging by using Amazon Elasticsearch Service (Amazon ES), Logstash, and Kibana.
- C. Scale down the application to one larger EC2 instance where only one instance is recording logs.
- D. Install the unified Amazon CloudWatch agent on the EC2 instance.
- E. Configure the agent to push the application logs to CloudWatch.

Answer: D

NEW QUESTION 465

- (Exam Topic 3)

A company hosts a microservices application that uses Amazon API Gateway, AWS Lambda, Amazon Simple Queue Service (Amazon SQS), and Amazon DynamoDB. One of the Lambda functions adds messages to an SQS FIFO queue.

When a developer checks the application logs, the developer finds a few duplicated items in a DynamoDB table. The items were inserted by another polling function that processes messages from the queue.

What is the MOST likely cause of this issue?

- A. Write operations on the DynamoDB table are being throttled.
- B. The SQS queue delivered the message to the function more than once.
- C. API Gateway duplicated the message in the SQS queue.
- D. The polling function timeout is greater than the queue visibility timeout.

Answer: B

NEW QUESTION 467

- (Exam Topic 3)

A developer used the BatchWriteItem API operation to insert items in an Amazon DynamoDB table. DynamoDB returned a few items as unprocessed due to

throttling The developer decides to retry the records on the unprocessed items
What should the developer do to reprocess the records with the LEAST number of API calls'?

- A. Retry the BatchWriteItem operation immediately
- B. Perform the PutItem operation on the unprocessed items individually instead of using the BatchWriteItem operation
- C. Delay the BatchWriteItem operation by using progressively longer wait times between retries, or exponential backoff
- D. Delete the items that were successfully processed, and reissue a new BatchWriteItem operation

Answer: D

NEW QUESTION 470

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-Developer-Associate Exam with Our Prep Materials Via below:

<https://www.certleader.com/AWS-Certified-Developer-Associate-dumps.html>