



Amazon

Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate

About ExamBible

Your Partner of IT Exam

Found in 1998

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

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

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

NEW QUESTION 1

- (Exam Topic 1)

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

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

Answer: C

NEW QUESTION 2

- (Exam Topic 1)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 3

- (Exam Topic 1)

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-frequency/>

NEW QUESTION 4

- (Exam Topic 1)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 5

- (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 6

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 7

- (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 8

- (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 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)

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 10

- (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 14

- (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 17

- (Exam Topic 1)

A Developer wants to enable AWS X-Ray for a secure application that runs in an Amazon ECS environment. What combination of steps will enable X-Ray? (Select THREE.)

- A. Create a Docker image that runs the X-Ray daemon.
- B. Add instrumentation to the application code for X-Ray.
- C. Install the X-Ray daemon on the underlying EC2 instance.
- D. Configure and use an IAM EC2 instance role.
- E. Register the application with X-Ray.
- F. Configure and use an IAM role for tasks.

Answer: ABF

Explanation:

<https://docs.aws.amazon.com/xray/latest/devguide/xray-daemon-ecs.html> <https://docs.aws.amazon.com/xray/latest/devguide/scorekeep-ecs.html>

NEW QUESTION 21

- (Exam Topic 1)

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

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

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

Answer: D

NEW QUESTION 22

- (Exam Topic 1)

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

What can the Developer do to fix this problem?

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

Answer: C

NEW QUESTION 23

- (Exam Topic 1)

A Developer is creating a serverless website with content that includes HTML files, images, videos, and JavaScript (client-side scripts). Which combination of services should the Developer use to create the website?

- A. Amazon S3 and Amazon CloudFront
- B. Amazon EC2 and Amazon ElastiCache
- C. Amazon ECS and Redis
- D. AWS Lambda and Amazon API Gateway

Answer: A

Explanation:

Reference: <https://d1.awsstatic.com/whitepapers/Building%20Static%20Websites%20on%20AWS.pdf>

NEW QUESTION 28

- (Exam Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 30

- (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 32

- (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 35

- (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 40

- (Exam Topic 1)

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

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

Answer: AC

Explanation:

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

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

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

“HTTP”, “HTTPS” – Subscribers specify a URL as part of the subscription registration; notifications will be delivered through an HTTP POST to the specified URL.

“Email”, “Email-JSON” – Messages are sent to registered addresses as email. Email-JSON sends notifications as a JSON object, while Email sends text-based email.

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

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

NEW QUESTION 44

- (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 49

- (Exam Topic 1)

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

In which order will these records be processed?

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

Answer: C

NEW QUESTION 53

- (Exam Topic 1)

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

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

- A. Use an Amazon CloudWatch event backed by an AWS Lambda function
- B. Issue an Amazon S3 API call to get a list of all Amazon S3 objects and persist the metadata within DynamoDB

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

Answer: A

NEW QUESTION 58

- (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%20task,IAM,role)

NEW QUESTION 63

- (Exam Topic 1)

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

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

Answer: CE

NEW QUESTION 64

- (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 65

- (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 67

- (Exam Topic 1)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 72

- (Exam Topic 1)

A company is creating an application that will require users to access AWS services and allow them to reset their own passwords. Which of the following would allow the company to manage users and authorization while allowing users to reset their own passwords?

- A. Amazon Cognito identify pools and AWS STS
- B. Amazon Cognito identity pools and AWS IAM
- C. Amazon Cognito user pools and AWS KMS
- D. Amazon Cognito user pools and identity pools

Answer: D

Explanation:

<https://serverless-stack.com/chapters/cognito-user-pool-vs-identity-pool.html>

NEW QUESTION 73

- (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 74

- (Exam Topic 1)

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

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

Answer: AC

Explanation:

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

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

NEW QUESTION 75

- (Exam Topic 1)

In AWS, which security aspects are the customer's responsibility? Choose 4 answers

- A. Life-cycle management of IAM credentials
- B. Decommissioning storage devices
- C. Security Group and ACL (Access Control List) settings
- D. Encryption of EBS (Elastic Block Storage) volumes
- E. Controlling physical access to compute resources
- F. Patch management on the EC2 instance's operating system

Answer: ACDF

Explanation:

Physical and Environmental Security

AWS's data centers are state of the art, utilizing innovative architectural and engineering approaches. Amazon has many years of experience in designing, constructing, and operating large-scale data centers. This experience has been applied to the AWS platform and infrastructure. AWS data centers are housed in nondescript facilities. Physical access is strictly controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff.

Storage Decommissioning

- When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals.
- AWS uses the techniques detailed in DoD 5220.22-M (National Industrial Security Program Operating Manual) or NIST 800-88 (Guidelines for Media Sanitization) to destroy data as part of the decommissioning process.
- All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

NEW QUESTION 76

- (Exam Topic 1)

A Developer is receiving HTTP 400: ThrottlingException errors intermittently when calling the Amazon CloudWatch API. When a call fails, no data is retrieved. What best practice should first be applied to address this issue?

- A. Contact AWS Support for a limit increase.
- B. Use the AWS CLI to get the metrics
- C. Analyze the applications and remove the API call
- D. Retry the call with exponential backoff

Answer: D

Explanation:

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cloudwatch_limits.html

NEW QUESTION 78

- (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 81

- (Exam Topic 1)

A Developer is creating a mobile application that will not require users to log in. What is the MOST efficient method to grant users access to AWS resources?

- A. Use an identity provider to securely authenticate with the application.
- B. Create an AWS Lambda function to create an IAM user when a user accesses the application.
- C. Create credentials using AWS KMS and apply these credentials to users when using the application.
- D. Use Amazon Cognito to associate unauthenticated users with an IAM role that has limited access to resources.

Answer: D

Explanation:

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

NEW QUESTION 86

- (Exam Topic 1)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 88

- (Exam Topic 1)

A Developer has created a software package to be deployed on multiple EC2 instances using IAM roles. What actions could be performed to verify IAM access to get records from Amazon Kinesis Streams? (Select TWO.)

- A. Use the AWS CLI to retrieve the IAM group.
- B. Query Amazon EC2 metadata for in-line IAM policies.
- C. Request a token from AWS STS, and perform a describe action.
- D. Perform a get action using the `--dry-run` argument.
- E. Validate the IAM role policy with the IAM policy simulator.

Answer: AE

NEW QUESTION 93

- (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 94

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 98

- (Exam Topic 1)

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

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

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

Answer: D

NEW QUESTION 99

- (Exam Topic 1)

Which of the following is chosen as the default region when making an API call with an AWS SDK?

- A. ap-northeast-1
- B. us-west-2
- C. us-east-1
- D. eu-west-1
- E. us-central-1

Answer: C

Explanation:

<https://docs.aws.amazon.com/sdk-for-java/v1/developer-guide/java-dg-region-selection.html>. This section applies only when using a client builder to access AWS services. AWS clients created by using the client constructor will not automatically determine region from the environment and will, instead, use the default SDK region (USEast1).

NEW QUESTION 103

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 108

- (Exam Topic 1)

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

```
A client error (UnauthorizedOperation) occurred when calling the RunInstances operation:
You are not authorized to perform this operation. Encoded authorization failure message:
oGsbAaIV7wlfj8zUqebHUANHzFbmkzILlxyj_y9xwhIhk99U_cUq1FIeZnskWDjQ1wSHStVfdCEyZILGoccGpCiC
IhORceWF9rRwFTnEcRj3N9iTrPAE1WHveC5Z54ALPaWlEjHlLg8CaB8d8lCKmxQuylCm0r1Bf2fHJRujAYopMVmga
8olFmKAl9yn_25rI120Q9p5ZIMX28zYM4dTulcJQUQjosgrEejfiIMYDda8l7Ooko9H6VmGJX62KfkRa5l7yE6hhh
2bIwA6tpyCJy2LWFRTe4bafqAyoqkarhPA4mGiZyWn4gSqbO8o-
uqSIvWYPweaKGkampa0arcFR4gBD7Ph097WYBkzX9hVjGppLMy4jpXRv
```

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 110

- (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 115

- (Exam Topic 1)

An application stores payroll information nightly in DynamoDB for a large number of employees across hundreds of offices. Item attributes consist of individual name, office identifier, and cumulative daily hours. Managers run reports for ranges of names working in their office. One query is, "Return all items in this office for names starting with A through E".

Which table configuration will result in the lowest impact on provisioned throughput for this query?

- A. Configure the table to have a hash index on the name attribute, and a range index on the office identifier
- B. Configure the table to have a range index on the name attribute, and a hash index on the office identifier
- C. Configure a hash index on the name attribute and no range index
- D. Configure a hash index on the office identifier attribute and no range index

Answer: B

Explanation:

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

Partition key and sort key – Referred to as a composite primary key, this type of key is composed of two attributes. The first attribute is the partition key, and the second attribute is the sort key. DynamoDB uses the partition key value as input to an internal hash function. The output from the hash function determines the partition (physical storage internal to DynamoDB) in which the item will be stored. All items with the same partition key value are stored together, in sorted order by sort key value.

NEW QUESTION 118

- (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 119

- (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 identifyproviders 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 ST
- 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 124

- (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 128

- (Exam Topic 1)

You run an ad-supported photo sharing website using S3 to serve photos to visitors of your site. At some point you find out that other sites have been linking to the photos on your site, causing loss to your business.

What is an effective method to mitigate this?

- A. Store photos on an EBS volume of the web server
- B. Remove public read access and use signed URLs with expiry dates.
- C. Use CloudFront distributions for static content.
- D. Block the IPs of the offending websites in Security Groups.

Answer: B

Explanation:

<https://aws.amazon.com/getting-started/projects/building-fast-session-caching-with-amazon-elasticache-for-redis/>

NEW QUESTION 133

- (Exam Topic 1)

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

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

- A. Ask customers to use AWS credentials to call the InvalidateCache API.
- B. Ask customers to invoke an AWS API endpoint which invalidates the cache.
- C. Ask customers to pass an HTTP header called Cache-Control:max-age=0.
- D. Ask customers to add a query string parameter called "INVALIDATE_CACHE" when making an API call.

Answer: C

Explanation:

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

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

NEW QUESTION 135

- (Exam Topic 1)

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

- A. IAM temporary credentials using AWS STS.

- B. Amazon Directory Service
- C. Amazon Cognito with unauthenticated access enabled
- D. IAM with SAML integration

Answer: C

Explanation:

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

NEW QUESTION 139

- (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 140

- (Exam Topic 1)

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

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

What is the MOST likely cause of this situation?

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

Answer: B

Explanation:

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

NEW QUESTION 144

- (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 147

- (Exam Topic 1)

What happens, by default, when one of the resources in a CloudFormation stack cannot be created?

- A. Previously-created resources are kept but the stack creation terminates.
- B. Previously-created resources are deleted and the stack creation terminates.
- C. The stack creation continues, and the final results indicate which steps failed.
- D. CloudFormation templates are parsed in advance so stack creation is guaranteed to succeed.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html>

AWS CloudFormation ensures all stack resources are created or deleted as appropriate. Because AWS CloudFormation treats the stack resources as a single unit, they must all be created or deleted successfully for the stack to be created or deleted. If a resource cannot be created, AWS CloudFormation rolls the stack back and automatically deletes any resources that were created. If a resource cannot be deleted, any remaining resources are retained until the stack can be successfully deleted.

NEW QUESTION 148

- (Exam Topic 2)

A Developer is writing an application that runs on Amazon EC2 instances in an Auto scaling group. The application data is stored in an Amazon DynamoDB table and records are constantly updated by all instances. An instance sometimes retrieves old data. The Developer wants to correct this by making sure the reads are strongly consistent.

How can the developer accomplish this?

- A. Set consistentRead to true when calling Getitem.
- B. Create a new DynamoDB Accelerator (DAX) table.
- C. Set consistency to strong when calling Update Table.
- D. Use the Getshardlterator command.

Answer: B

Explanation:

Reference: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadConsistency.html>

NEW QUESTION 152

- (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 154

- (Exam Topic 2)

A company wants to containerize an existing three-tier web application and deploy it to Amazon ECS Fargate. The application is using session data to keep track of user activities.

Which approach would provide the BEST user experience?

- A. Provision a Redis cluster in Amazon ElasticCache and save the session data in the cluster
- B. Create a session table in Amazon Redshift and save the session data in the database table.
- C. Enable session stickiness in the existing Network Load Balancer and manage the session data in the container.
- D. Use an Amazon S3 bucket as data store and save the session data in the bucket.

Answer: C

NEW QUESTION 156

- (Exam Topic 2)

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

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

Answer: AC

NEW QUESTION 158

- (Exam Topic 2)

A developer is preparing a deployment package using AWS Cloud Formation. The package consists of two separate templates: one for the infrastructure and one for the application. The application has to be inside the VPC that is created from the infrastructure template

How can the application stack refer to the VPC created from the infrastructure template?

- A. Use the Ret function to import the VPC into the application stack from the infrastructure template
- B. Use the export flag in the infrastructure template, and then use the Fn::ImportValue function in the application template
- C. Use the DependsOn attribute to specify that the application instance depends on the VPC in the application template
- D. Use the Fn::GetAtt function to include the attribute of the VPC in the application template.

Answer: A

NEW QUESTION 162

- (Exam Topic 2)

An e-commerce web application that shares session state on-premises is being migrated to AWS. The application must be fault tolerant, natively highly scalable, and any service interruption should not affect the user experience.

What is the best option to store the session state?

- A. Store the session state in Amazon ElastiCache
- B. Store the session state in Amazon CloudFront

- C. Store the session state in Amazon S3
- D. Enable session stickiness using elastic load balancers

Answer: A

Explanation:

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

NEW QUESTION 167

- (Exam Topic 2)

An application is running on a cluster of Amazon EC2 instance. While trying to read objects stored within a single Amazon S3 bucket that are encrypted with server-side encryption with AWS KMS managed keys (SSE-KMS), the application receives the following error:

Service : AWSKMS: Status Code: 400: Code : ThrottlingException

Which combination of steps should be taken to prevent this failure? (Select TWO.)

- A. Contact AWS Support to request an AWS KMS rate limit increase.
- B. Perform error retries with exponential backoff in the application code.
- C. Contact AWS Support to request a S3 rate limit increase.
- D. Import a customer master key (CMK) with a larger key size.
- E. Use more than one customer master key (CMK) to encrypt S3 data

Answer: AD

NEW QUESTION 170

- (Exam Topic 2)

A company is creating a REST service using an Amazon API Gateway with AWS Lambda integration. The service must run different versions for testing purposes. What would be the BEST way to accomplish this?

- A. Use an X-Version header to denote which version is being called and pass that header to the Lambda function(s)
- B. Create an API Gateway Lambda authorizer to route API clients to the correct API version
- C. Create an API Gateway resource policy to isolate versions and provide context to the Lambda function(s)
- D. Deploy the API versions as unique stages with unique endpoints and use stage variables to provide further context

Answer: A

Explanation:

The service run different versions for testing purposes. means different stages, stage variable is the way to go.

NEW QUESTION 171

- (Exam Topic 2)

A company is adding stored value for gift card) capability to its highly popular casual gaming website. Users need to be able to trade this value for other users' items on the platform. This would require both users' records be updated as a single transaction, or both users' records to be completely rolled back.

Which AWS database options can provide the transactional capability required for this new feature? (Select TWO)

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

Answer: CD

NEW QUESTION 172

- (Exam Topic 2)

A company runs an e-commerce website that uses Amazon DynamoDB where pricing for items is dynamically updated in real time. At any given time, multiple updates may occur simultaneously for pricing information on a particular product. This is causing the original editor's changes to be overwritten without a proper review process.

Which DynamoDB write option should be selected to prevent this overwriting?

- A. Concurrent writes
- B. Conditional writes
- C. Atomic writes
- D. Batch writes

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/WorkingWithItems.html#WorkingWithIt>

NEW QUESTION 174

- (Exam Topic 2)

An AWS Lambda function must access an external site by using a regularly rotated user name and password. These items must be kept securely and cannot be stored in the function code.

What combination of AWS services can be used to accomplish this? (Choose two.)

- A. AWS Certificate Manager (ACM)
- B. AWS Systems Manager Parameter Store
- C. AWS Trusted Advisor

- D. AWS KMS
- E. Amazon GuardDuty

Answer: BD

Explanation:

<https://docs.aws.amazon.com/kms/latest/developerguide/services-parameter-store.html>

NEW QUESTION 178

- (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 182

- (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 185

- (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 190

- (Exam Topic 2)

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

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

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

Answer: D

NEW QUESTION 192

- (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 194

- (Exam Topic 2)

A developer registered an AWS Lambda function as a target for an Application Load Balancer (ALB) using a CLI command. However, the Lambda function is not being invoked when the client sends requests through the ALB.

Why is the Lambda function not being invoked?

- A. A Lambda function cannot be registered as a target for an ALB
- B. A Lambda function can be registered with an ALB using AWS Management Console only
- C. The permissions to invoke the Lambda function are missing
- D. Cross-zone is not enabled on the ALB

Answer: C

Explanation:

Reference: <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/lambda-functions.html>

NEW QUESTION 199

- (Exam Topic 2)

A developer wants to build an application that will allow new users to register and create new user accounts. The application must also allow users with social media accounts to log in using their social media credentials.

Which AWS service or feature can be used to meet these requirements?

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 201

- (Exam Topic 2)

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

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

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

Answer: B

NEW QUESTION 204

- (Exam Topic 2)

A developer needs to create an application that supports Security Assertion Markup Language (SAML) and Facebook authentication. It must also allow access to AWS services, such as Amazon DynamoDB.

Which AWS service or feature will meet these requirements with the LEAST amount of additional coding?

- A. AWSAppSync
- B. Amazon Cognito identity pools
- C. Amazon Cognito user pools
- D. Amazon Lambda@Edge

Answer: B

Explanation:

Reference: <https://aws.amazon.com/blogs/mobile/amazon-cognito-user-pools-supports-federation-with-saml/>

NEW QUESTION 209

- (Exam Topic 2)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 210

- (Exam Topic 2)

A developer is implementing authentication and authorization for an application. The developer needs to ensure that the user credentials are never exposed.

Which approach should the developer take to meet this requirement?

- A. Store the user credentials In Amazon DynamoDB Build an AWS Lambda function to validate the credentials and authorize users
- B. Deploy a custom authentication and authorization API on an Amazon EC2 instance
- C. Store the user credentials in Amazon S3 and encrypt the credentials using Amazon S3 server-side encryption.
- D. Use Amazon Cognito to configure a user pool, and use the Cognito API to authenticate and authorize the users.
- E. Store the user credentials In Amazon RDS Enable the encryption option for the Amazon RDS D8 instances Build an API using AWS Lambda to validate the credentials and authorize users

Answer: C

NEW QUESTION 211

- (Exam Topic 2)

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

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

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

Answer: CD

Explanation:

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

NEW QUESTION 212

- (Exam Topic 2)

A company has implemented AWS CodeDeploy as part of its cloud native CI/CD stack The company enables automatic rollbacks while deploying a new version of a popular web application from in place to Amazon EC2.

What occurs if the deployment of the new version fails due to code regression?

- A. The last known good deployment is automatically restored using the snapshot stored in Amazon S3
- B. CodeDeploy switches the Amazon Route 53 alias records back to the known good green deployment and terminates the failed blue deployment
- C. A new deployment of the last known good version of the application is deployed with a new deployment ID
- D. AWS CodePipeline promotes the most recent deployment with a SUCCEEDED status to production

Answer: B

NEW QUESTION 217

- (Exam Topic 2)

A developer uses Amazon S3 buckets for static website hosting. The developer creates one S3 bucket for the code and another S3 bucket for the assets, such as image and video files. Access is denied when a user attempts to access the assets bucket from the code bucket, with the website application showing a 403 error How should the developer solve this issue?

- A. Create an IAM role and apply it to the assets bucket for the code bucket to be granted access
- B. Edit the bucket policy of the assets bucket to open access to all principals
- C. Edit the cross-origin resource sharing (CORS) configuration of the assets bucket to allow any origin to access the assets
- D. Change the code bucket to use AWS Lambda functions instead of static website hosting.

Answer: C

NEW QUESTION 220

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

https://docs.aws.amazon.com/en_us/amazondynamodb/latest/developerguide/Streams.Lambda.Tutorial.html Create an event source mapping to tell Lambda to send records from your stream to a Lambda function. You can create multiple event source mappings to process the same data with multiple Lambda functions, or process items from multiple streams with a single function.

NEW QUESTION 224

- (Exam Topic 2)

A company runs continuous integration/continuous delivery (CI/CD) pipeline for its application on AWS CodePipeline. A developer must write unit tests and run them as part of the pipelines before staging the artifacts for testing.

How should the Developer incorporate unit tests as part of CI/CD pipeline?

- A. Create a separate codePipeline pipeline to run unit tests.
- B. Update the AWS codeBuild build specification to include a phase for running unit tests.
- C. Install the AWS CodeDeploy agent on an Amazon EC2 instance to run unit tests.
- D. Create a testing branch in AWS CodeCommit to run unit tests.

Answer: B

NEW QUESTION 228

- (Exam Topic 2)

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

How can a developer meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 231

- (Exam Topic 2)

An organization is using Amazon CloudFront to ensure that its users experience low-latency access to its web application. The organization has identified a need to encrypt all traffic between users and CloudFront, and all traffic between CloudFront and the web application.

How can these requirements be met? (Choose two.)

- A. Use AWS KMS to encrypt traffic between CloudFront and the web application.
- B. Set the Origin Protocol Policy to "HTTPS Only".
- C. Set the Origin's HTTP Port to 443.
- D. Set the Viewer Protocol Policy to "HTTPS Only" or "Redirect HTTP to HTTPS".
- E. Enable the CloudFront option Restrict Viewer Access.

Answer: AB

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using-https-viewers-to-cloudfront.htm>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using-https-cloudfront-to-custom-origi>

NEW QUESTION 233

- (Exam Topic 2)

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

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

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

Answer: BD

Explanation:

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

NEW QUESTION 235

- (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 240

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 245

- (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 248

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 253

- (Exam Topic 2)

A developer is creating a new application that will be accessed by users through an API created using Amazon API Gateway. The users need to be authenticated by a third-party Security Assertion Markup Language (SAML) identity provider. Once authenticated, users will need access to other AWS services such as Amazon S3 and Amazon DynamoDB. How can these requirements be met?

- A. Use an Amazon Cognito user pool with SAML as the resource server
- B. Use Amazon Cognito Identity pools with a SAML identity provider as one of the authentication providers
- C. Use the AWS IAM service to provide the sign-up and sign-in functionality.
- D. Use Amazon CloudFront signed URLs to connect with the SAML identity provider

Answer: B

NEW QUESTION 257

- (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 258

- (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 259

- (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 262

- (Exam Topic 2)

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

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

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

Answer: C

NEW QUESTION 267

- (Exam Topic 2)

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

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

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

Answer: DE

NEW QUESTION 268

- (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 273

- (Exam Topic 2)

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

Error: aws: command not found

What is the most likely cause of this error?

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

Answer: A

Explanation:

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

NEW QUESTION 276

- (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 277

- (Exam Topic 2)

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

How can the Developer MOST efficiently handle the temporary files?

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

Answer: C

NEW QUESTION 281

- (Exam Topic 2)

A Developer must analyze performance issues with production-distributed applications written as AWS Lambda functions. These distributed Lambda applications invoke other components that make up the applications.

How should the Developer identify and troubleshoot the root cause of the performance issues in production?

- A. Add logging statements to the Lambda functions, then use Amazon CloudWatch to view the logs.
- B. Use AWS Cloud Trail and then examine the logs
- C. Use AWS X-Ray, then examine the segments and errors
- D. Run Amazon Inspector agents and then analyze performance

Answer: C

Explanation:

<https://aws.amazon.com/blogs/developer/new-analyze-and-debug-distributed-applications-interactively-using-aw>

NEW QUESTION 283

- (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 288

- (Exam Topic 2)

A developer wants to ensure the Amazon EC2 instances in AWS Elastic Beanstalk execute a certain set of commands before the application is ready to use Which Elastic Beanstalk feature will allow the developer to accomplish this?

- A. Rolling update
- B. Immutable update
- C. User data
- D. ebextensions

Answer: D

NEW QUESTION 293

- (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 296

- (Exam Topic 2)

A Developer is migrating existing applications to AWS. These applications use MongoDB as their primary data store, and they will be deployed to Amazon EC2 instances. Management requires that the Developer minimize changes to applications while using AWS services Which solution should the Developer use to host MongoDB in AWS?

- A. Install MongoDB on the same instance where the application is running
- B. Deploy Amazon DocumentDB in MongoDB compatibility mode
- C. Use Amazon API Gateway to translate API calls from MongoDB to Amazon DynamoDB.
- D. Replicate the existing MongoDB workload to Amazon DynamoDB

Answer: D

NEW QUESTION 298

- (Exam Topic 2)

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

What should the development team do to meet these requirements?

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

Answer: C

NEW QUESTION 299

- (Exam Topic 2)

A company needs a new REST API that can return information about the contents of an Amazon S3 bucket, such as a count of the objects stored in it. The company has decided that the new API should be written as a microservice using AWS Lambda and Amazon API Gateway.

How should the Developer ensure that the microservice has the necessary access to the Amazon S3 bucket, while adhering to security best practices?

- A. Create an IAM user that has permissions to access the Amazon S3 bucket, and store the IAM user credentials in the Lambda function source code.
- B. Create an IAM role that has permissions to access the Amazon S3 bucket and assign it to the Lambda function as its execution role.
- C. Create an Amazon S3 bucket policy that specifies the Lambda service as its principal and assign it to the Amazon S3 bucket.
- D. Create an IAM role, attach the AmazonS3FullAccess managed policy to it, and assign the role to the Lambda function as its execution role.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/lambda-execution-role-s3-bucket/>

NEW QUESTION 300

- (Exam Topic 2)

A stock market monitoring application uses Amazon Kinesis for data ingestion. During simulated tests of peak data rates, the Kinesis stream cannot keep up with the incoming data.

What step will allow Kinesis to accommodate the traffic during peak hours?

- A. Install the Kinesis Producer Library (KPL) for ingesting data into the stream.
- B. Reduce the data retention period to allow for more data ingestion using `DecreaseStreamRetentionPeriod`.
- C. Increase the shard count of the stream using `UpdateShardCount`.
- D. Ingest multiple records into the stream in a single call using `PutRecords`.

Answer: C

Explanation:

<https://docs.aws.amazon.com/streams/latest/dev/developing-producers-with-kpl.html>

NEW QUESTION 303

- (Exam Topic 2)

A company is migrating from a monolithic architecture to a microservices-based architecture. The Developers need to refactor the application so that the many microservices can asynchronously communicate with each other without impacting performance.

Use of which managed AWS services will enable asynchronous message passing? (Choose two.)

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

Answer: AD

NEW QUESTION 306

- (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 310

- (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 repositor
- 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 314

- (Exam Topic 2)

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

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 317

- (Exam Topic 2)

A company experienced partial downtime during the last deployment of a new application AWS Elastic Beanstalk split the environment's Amazon EC2 instances into batches and deployed a new version one batch at a time after taking them out of service. Therefore, full capacity was not maintained during deployment. The developer plans to release a new version of the application, and is looking for a policy that will maintain full capacity and minimize the impact of the failed deployment

Which deployment policy should the developer use?

- A. Immutable
- B. All at Once
- C. Rolling
- D. Rolling with an Additional Batch

Answer: A

Explanation:

Immutable infrastructure has become a new norm in IT operations. Immutable Deployment is one of those approaches, and it simply means: Immutable: the "staging" environment, once ready to become production, doesn't change. If we need to change something, we then deploy new code on completely new infrastructure. The benefits of an immutable infrastructure include more consistency and reliability in your infrastructure and a simpler, more predictable deployment process

NEW QUESTION 318

- (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 323

- (Exam Topic 3)

A developer is building an application that processes a stream of user-supplied data The data stream must be consumed by multiple Amazon EC2 based processing applications in parallel and in real time. Each processor must be able to resume without losing data if there is a service interruption. The Application Architect plans to add other processors in the near future, and wants to minimize the amount of data duplication involved.

Which solution will satisfy these requirements?

- A. Publish the data to Amazon SQS
- B. Publish the data to Amazon Kinesis Data Firehose
- C. Publish the data to Amazon CloudWatch Events.
- D. Publish the data to Amazon Kinesis Data Streams.

Answer: A

NEW QUESTION 324

- (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 326

- (Exam Topic 3)

Which of the following are good use cases for how Amazon ElastiCache can help an application? (Select TWO.)

- A. Improve the performance of S3 PUT operations
- B. Improve the latency of deployments performed by AWS CodeDeploy
- C. Improve latency and throughput for read-heavy application workloads.
- D. Reduce the time required to merge AWS CodeCommit branches
- E. Improve performance of compute-intensive applications.

Answer: CE

NEW QUESTION 329

- (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 330

- (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 333

- (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 338

- (Exam Topic 3)

A developer wants to use React to build a web and mobile application. The application will be hosted on AWS The application must authenticate users and then allow users to store and retrieve files that they own The developer wants to use Facebook for authentication

Which CLI will MOST accelerate the development and deployment of this application on AWS?

- A. AWS CLI
- B. AWS Amplify CLI
- C. AWS Serverless Application Model (AWS SAM) CLI
- D. Amazon Elastic Container Service (Amazon ECS) CLI

Answer: B

NEW QUESTION 339

- (Exam Topic 3)

A developer must allow guest users without logins to access an Amazon Cognito-enabled site to view files stored within an Amazon S3 bucket How should the developer meet these requirements'?

- A. Create a blank user ID in a user pool, add to the user group, and grant access to AWS resources
- B. Create a new identity pool, enable access to unauthenticated identities and grant access to AWS resources
- C. Create a new user pool, enable access to unauthenticated identities, and grant access to AWS resources.

D. Create a new user pool disable authentication access, and grant access to AWS resources

Answer: C

NEW QUESTION 343

- (Exam Topic 3)

A developer must extend an existing application that is based on the AWS Services Application Model (AWS SAM). The developer has used the AWS SAM CLI to create the project. The project contains different AWS Lambda functions.

Which combination of commands must the developer use to redeploy the AWS SAM application (Select TWO.)

- A. Sam init
- B. Sam validate
- C. Sam build
- D. Sam deploy
- E. Sam publish

Answer: AD

NEW QUESTION 347

- (Exam Topic 3)

A developer is building an application integrating an Amazon API Gateway with an AWS Lambda function. When calling the API, the developer receives the following error. Wed Nov 03 01:13:00 UTC 2017 : Method completed with status: 502 What should the developer do to resolve the error?

- A. Change the HTTP endpoint of the API to an HTTPS endpoint.
- B. Change the format of the payload sent to the API Gateway.
- C. Change the format of the Lambda function response to the API call.
- D. Change the authorization header in the API call to access the Lambda function.

Answer: C

NEW QUESTION 350

- (Exam Topic 3)

A developer has written a multi-threaded application that is running on a fleet of Amazon EC2 instances. The operations team has requested a graphical method to monitor the number of running threads over time.

What is the MOST efficient way to fulfill this request?

- A. Periodically send the thread count to AWS X-Ray segments, then generate a service graph on demand
- B. Create a custom Amazon CloudWatch metric and periodically perform a PutMetricData call with the current thread count.
- C. Periodically log thread count data to Amazon S3. Use Amazon Kinesis to process the data into a graph.
- D. Periodically write the current thread count to a table using Amazon DynamoDB and use Amazon CloudFront to create a graph

Answer: D

NEW QUESTION 351

- (Exam Topic 3)

A developer is attempting to use the Amazon S3 PutObject API operation to upload an object to an S3 bucket that has default encryption enabled. The developer receives a 400 Bad Request error.

What is the MOST likely cause of this error?

- A. The API operation cannot access the encryption key
- B. The HTTP Content-Length header is missing.
- C. The object exceeds the maximum object size that is allowed.
- D. The S3 bucket exceeds the maximum storage capacity that is allowed

Answer: D

NEW QUESTION 354

- (Exam Topic 3)

A company's ecommerce 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 358

- (Exam Topic 3)

A company has three AWS Lambda functions that are written in Node.js. The Lambda functions include a mix of custom code and open-source modules. When bugs are occasionally detected in the open-source modules, all three Lambda functions must be patched.

What is the MOST operationally efficient solution to deploy a patched open-source library for all three Lambda functions?

- A. Create a custom AWS CloudFormation public registry extension Reference a GitHub repository that hosts the open-source modules in the extension Configure CloudFormation to scan the repository once each day Write an AWS Serverless Application Model (AWS SAM) template to redeploy the three Lambda functions upon a scan notification change.
- B. Create an Amazon CloudFront distribution with an Amazon S3 bucket as the origin Upload the patched modules to Amazon S3 when needed Modify each Lambda function to download the patched modules from the CloudFront distribution during the cold start.
- C. Launch an Amazon EC2 instance Host a private open-source module registry on the EC2 instance Upload the modified open-source modules to the private registry when needed
- D. Modify each Lambda function deployment script to download the modules from the private registry Redeploy the three new Lambda functions.
- E. Create a Lambda layer with the open-source modules Modify all three Lambda functions to depend on the layer Remove the open-source modules from each Lambda function Patch the Lambda layer with the modified open-source modules when needed Update the Lambda functions to reference the new layer version

Answer: D

NEW QUESTION 362

- (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 367

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 372

- (Exam Topic 3)

A developer needs to modify an application architecture to meet new functional requirements. Application data is stored in Amazon DynamoDB and processed for analysis in a nightly batch. The system analysts do not want to wait until the next day to view the processed data and have asked to have it available in near-real time.

Which application architect pattern would enable the data to be processed as it is received?

- A. Event driven
- B. Client served driven
- C. Fan-out driven
- D. Schedule driven

Answer: A

NEW QUESTION 373

- (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 378

- (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 380

- (Exam Topic 3)

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

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

How should a developer modify the application to improve performance?

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

Answer: A

NEW QUESTION 381

- (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 ID.
- D. The DynamoDB table is not running in the same Availability Zone as the Lambda function.

Answer: C

NEW QUESTION 384

- (Exam Topic 3)

A developer is creating AWS CloudFormation templates to manage an application's deployment in Amazon Elastic Container Service (Amazon ECS) through AWS CodeDeploy. The developer wants to automatically deploy new versions of the application to a percentage of users before the new version becomes available for all users.

How should the developer manage the deployment of the new version?

- A. Modify the CloudFormation template to include a Transform section and the AWS "CodeDeploy::BlueGreen" hook.
- B. Deploy the new version in a new CloudFormation stack. After testing is complete, update the application's DNS records for the new stack.
- C. Run CloudFormation stack updates on the application stack to deploy new application versions when they are available.
- D. Create a nested stack for the new version.
- E. Include a Transform section and the AWS: CodeDeploy BlueGreen hook.

Answer: B

NEW QUESTION 386

- (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 389

- (Exam Topic 3)

A development team decides to adopt a continuous integration/continuous delivery (CI/CD) process using AWS CodePipeline and AWS CodeCommit for a new application. However, management wants a person to review and approve the code before it is deployed to production.

How can the development team add a manual approver to the CI/CD pipeline?

- A. Use AWS SES to send an email to approvers when their action is required. Develop a simple application that allows approvers to accept or reject a build. Invoke an AWS Lambda function to advance the pipeline when a build is accepted.
- B. If approved, add an approved tag when pushing changes to the CodeCommit repository.
- C. CodePipeline will proceed to build and deploy approved commits without interruption.
- D. Add an approval step to CodeCommit. Commits will not be saved until approved.
- E. Add an approval action to the pipeline.
- F. Configure the approval action to publish to an Amazon SNS topic when approval is required.
- G. The pipeline execution will stop and wait for an approval.

Answer: D

NEW QUESTION 393

- (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 398

- (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 400

- (Exam Topic 3)

A company is launching a polling application. The application will store the results of each poll in an Amazon DynamoDB table. Management wants to remove poll data after a few days and store an archive of those records in Amazon S3.

Which approach would allow the application to archive each poll's data while keeping complexity to a MINIMUM?

- A. Enable Time to Live (TTL) on the DynamoDB table
- B. Enable DynamoDB Streams on the table and store the records removed from the stream in Amazon S3.
- C. Schedule an AWS Lambda function to periodically scan the DynamoDB table
- D. Use the BatchWriteItem operation to delete the results of a scan. Enable DynamoDB Stream on the table and store the records removed from the stream in Amazon S3.
- E. Enable DynamoDB Streams on the table
- F. Configure the stream as trigger for AWS Lambda
- G. Save records to Amazon S3 when records on the stream are modified.
- H. Enable cross-Region replication on the S3 bucket to archive the poll data.

Answer: C

NEW QUESTION 402

- (Exam Topic 3)

A company's fleet of Amazon EC2 instances receives data from millions of users through an API. The servers batch the data, add an object for each user, and upload the objects to an S3 bucket to ensure high access rates. The object attributes are Customer ID, Server ID, TS-Server (TimeStamp and Server ID), the size of the object, and a timestamp. A developer wants to find all the objects for a given user collected during a specified time range.

After creating an S3 object creation event, how can the developer achieve this requirement?

- A. Execute an AWS Lambda function in response to the S3 object creation events that creates an Amazon DynamoDB record for every object with the Customer ID as the partition key and the Server ID as the sort key. Retrieve all the records using the Customer ID and Server ID attributes.
- B. Execute an AWS Lambda function in response to the S3 object creation events that creates an Amazon Redshift record for every object with the Customer ID as the partition key and TS-Server as the sort key. Retrieve all the records using the Customer ID and TS-Server attributes.
- C. Execute an AWS Lambda function in response to the S3 object creation events that creates an Amazon DynamoDB record for every object with the Customer ID as the partition key and TS-Server as the sort key. Retrieve all the records using the Customer ID and TS-Server attributes.
- D. Execute an AWS Lambda function in response to the S3 object creation events that creates an Amazon Redshift record for every object with the Customer ID as the partition key and the Server ID as the sort key.
- E. Retrieve all the records using the Customer ID and Server ID attributes.

Answer: C

NEW QUESTION 404

- (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 406

- (Exam Topic 3)

A developer is building an application that runs behind an application Load Balancer (ALB). The application is configured as the origin for an Amazon CloudFront distribution. Users will log in to the application using their social media accounts.

How can the developer authenticate and authorize users?

- A. Validate the user by inspecting the tokens using AWS Lambda authorizers on the ALB
- B. Configure the ALB to use Amazon Cognito as one of the authentication providers
- C. Configure Cloudfront to use Amazon Cognito as one of the authentication providers
- D. Authorize the users by calling the Amazon Cognito API in the AWS Lambda authorizer on the ALB

Answer: C

NEW QUESTION 408

- (Exam Topic 3)

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

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

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

Answer: DE

NEW QUESTION 409

- (Exam Topic 3)

A company has an internal website that gives users the ability to access contract data that is stored in an Amazon RDS DB instance. The number of contracts has increased and several users have reported slow retrieval of the contract data.

The company wants to set up a cache to improve the latency. A developer must create a solution that ensures data resiliency. The data must be encrypted and must be partitioned by department.

Which solution will meet these requirements?

- A. Amazon ElastiCache for Memcached with cluster mode enabled
- B. Amazon ElastiCache for Redis with cluster mode enabled
- C. Amazon ElastiCache for Redis with cluster mode disabled
- D. Amazon ElastiCache for Memcached with cluster mode disabled

Answer: C

NEW QUESTION 411

- (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 414

- (Exam Topic 3)

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

Where should the temporary file be stored?

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

Answer: A

NEW QUESTION 419

- (Exam Topic 3)

A developer is writing an application to analyze the traffic to a fleet of Amazon EC2 instances. The EC2 instances run behind a public Application Load Balancer (ALB). An HTTP server runs on each of the EC2 instances, logging all requests to a log file.

The developer wants to capture the client public IP addresses. The developer analyzes the log files and notices only the IP address of the ALB.

What must the developer do to capture the client public IP addresses in the log file?

- A. Add a Host header to the HTTP server log configuration file
- B. Install the Amazon CloudWatch Logs agent on each EC2 instance
- C. Configure the agent to write to the log file.
- D. Install the AWS X-Ray daemon on each EC2 instance Configure the daemon to write to the log file.
- E. Add an X-Forwarded-For header to the HTTP server log configuration file.

Answer: C

NEW QUESTION 421

.....

Relate Links

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

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

Contact us

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

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