

AWS-Certified-Developer-Associate Dumps

Amazon AWS Certified Developer - Associate

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



NEW QUESTION 1

- (Exam Topic 1)

Your application is trying to upload a 6 GB file to Simple Storage Service and receive a "Your proposed upload exceeds the maximum allowed object size." error message.

What is a possible solution for this?

- A. None, Simple Storage Service objects are limited to 5 GB
- B. Use the multi-part upload API for this object
- C. Use the large object upload API for this object
- D. Contact support to increase your object size limit
- E. Upload to a different region

Answer: B

Explanation:

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

NEW QUESTION 2

- (Exam Topic 1)

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

What step is required to update the function with the changes?

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

Answer: B

Explanation:

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

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

NEW QUESTION 3

- (Exam Topic 1)

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



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

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

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

Answer: BD

Explanation:

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

NEW QUESTION 4

- (Exam Topic 1)

A Developer must re-implement the business logic for an order fulfilment 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 company has three different environments: Development, QA, and Production. The company wants to deploy its code first in the Development environment, then QA, and then Production.

Which AWS service can be used to meet this requirement?

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

Answer: D

Explanation:

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

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

NEW QUESTION 6

- (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 7

- (Exam Topic 1)

How can you secure data at rest on an EBS volume?

- A. Attach the volume to an instance using EC2's SSL interface.
- B. Write the data randomly instead of sequentially.
- C. Use an encrypted file system on top of the EBS volume.
- D. Encrypt the volume using the S3 server-side encryption service.
- E. Create an IAM policy that restricts read and write access to the volume.

Answer: C

NEW QUESTION 8

- (Exam Topic 1)

Which of the following platforms are supported by Elastic Beanstalk? Choose 2 answers

- A. Apache Tomcat
- B. .NET
- C. IBM Websphere
- D. Oracle JBoss
- E. Jetty

Answer: AB

Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/platforms/platforms-supported.html>

NEW QUESTION 9

- (Exam Topic 1) Company

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

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

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

Answer: AC

Explanation:

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

NEW QUESTION 10

- (Exam Topic 1)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 10

- (Exam Topic 1)

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

What is a possible solution to prevent this happening?

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

Answer: D

Explanation:

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

NEW QUESTION 15

- (Exam Topic 1)

A corporate web application is deployed within an Amazon VPC, and is connected to the corporate data center via IPsec VPN. The application must authenticate against the on-premise LDAP server. Once authenticated, logged-in users can only access an S3 keyspace specific to the user.

Which two approaches can satisfy the objectives? Choose 2 answers

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

Answer: BD

Explanation:

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

NEW QUESTION 17

- (Exam Topic 1)

A Developer uses AWS CodeDeploy to automate application deployment that connects to an external MySQL database. The Developer wants to securely access the encrypted secrets, such as API keys and database passwords.

Which of the following solutions would involve the LEAST administrative effort?

- A. Save the secrets in Amazon S3 with AWS KMS server-side encryption, and use a signed URL to access them by using the IAM role from Amazon EC2 instances.
- B. Use the instance metadata to store the secrets and to programmatically access the secrets from EC2 instances.
- C. Use the Amazon DynamoDB client-side encryption library to save the secrets in DynamoDB and to programmatically access the secrets from EC2 instances.
- D. Use AWS SSM Parameter Store to store the secrets and to programmatically access them by using the IAM role from EC2 instances.

Answer: D

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-parameter-store.html>

NEW QUESTION 20

- (Exam Topic 1)

A Developer wants to find a list of items in a global secondary index from an Amazon DynamoDB table. Which DynamoDB API call can the Developer use in order to consume the LEAST number of read capacity units?

- A. Scan operation using eventually-consistent reads
- B. Query operation using strongly-consistent reads
- C. Query operation using eventually-consistent reads
- D. Scan operation using strongly-consistent reads

Answer: C

Explanation:

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

NEW QUESTION 21

- (Exam Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 26

- (Exam Topic 1)

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

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

Answer: A

NEW QUESTION 29

- (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 30

- (Exam Topic 1)

Which of the following is an example of a good DynamoDB hash key schema for provisioned throughput efficiency?

- A. User ID, where the application has many different users.
- B. Status Code where most status codes are the same
- C. Device ID, where one is by far more popular than all the others.
- D. Game Type, where there are three possible game types

Answer: A

NEW QUESTION 35

- (Exam Topic 1)

A Developer has created an S3 bucket s3://mycoolapp and has enabled server access logging that points to the folder s3://mycoolapp/logs. The Developer moved 100 KB of Cascading Style Sheets (CSS) documents to the folder s3://mycoolapp/css, and then stopped work. When the developer came back a few days later, the bucket was 50 GB.

What is the MOST likely cause of this situation?

- A. The CSS files were not compressed and S3 versioning was enabled.
- B. S3 replication was enabled on the bucket.
- C. Logging into the same bucket caused exponential log growth.
- D. An S3 lifecycle policy has moved the entire CSS file to S3 Infrequent Access.

Answer: C

Explanation:

Refer AWS documentation - S3 Server logs

To turn on log delivery, you provide the following logging configuration information:

➤ The name of the target bucket where you want Amazon S3 to save the access logs as objects. You can have logs delivered to any bucket that you own that is in the same Region as the source bucket, including the source bucket itself. We recommend that you save access logs in a different bucket so that you can easily manage the logs. If you choose to save access logs in the source bucket, we recommend that you specify a prefix for all log object keys so that the object names begin with a common string and the log objects are easier to identify. When your source bucket and target bucket are the same bucket, additional logs are created for the logs that are written to the bucket. This behavior might not be ideal for your use case because it could result in a small increase in your storage billing. In addition, the extra logs about logs might make it harder to find the log that you're looking for.

NEW QUESTION 38

- (Exam Topic 1)

A company needs to secure its existing website running behind an Elastic Load Balancer. The website's Amazon EC2 instances are CPU-constrained.

What should be done to secure the website while not increasing the CPU load on the EC2 web servers? (Select TWO.)

- A. Configure an Elastic Load Balancer with SSL pass-through.
- B. Configure SSL certificates on an Elastic Load Balancer.
- C. Configure an Elastic Load Balancer with a Loadable Storage System.
- D. Install SSL certificates on the EC2 instances.
- E. Configure an Elastic Load Balancer with SSL termination.

Answer: BE

NEW QUESTION 43

- (Exam Topic 1)

A web application is using Amazon Kinesis Streams for clickstream data that may not be consumed for up to 12 hours.

How can the Developer implement encryption at rest for data within the Kinesis Streams?

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

Answer: D

Explanation:

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

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

NEW QUESTION 48

- (Exam Topic 1)

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

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

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

Answer: C

NEW QUESTION 53

- (Exam Topic 1)

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

How should the Developer identify the cause of the failures?

- A. Modify the buildspec.yml file to include steps to send the output of build commands to Amazon CloudWatch.
- B. Use a custom Docker image that includes the AWS X-Ray agent in the AWS CodeBuild project configuration.
- C. Check the build logs of the failed phase in the last build attempt in the AWS CodeBuild project build history.

D. Manually re-run the build process on a local machine so that the output can be visualized.

Answer: A

NEW QUESTION 55

- (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 59

- (Exam Topic 1)

A Developer is creating a mobile application with a limited budget. The solution requires a scalable service that will enable customers to sign up and authenticate into the mobile application while using the organization's current SAML 2.0 identity provider.

Which AWS service should be used to meet these requirements?

- A. AWS Lambda
- B. Amazon Cognito
- C. AWS IAM
- D. Amazon EC2

Answer: B

NEW QUESTION 61

- (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 66

- (Exam Topic 1)

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

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

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

Answer: C

NEW QUESTION 71

- (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 73

- (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 75

- (Exam Topic 1)

In a move toward using microservices, a company's Management team has asked all Development teams to build their services so that API requests depend only on that service's data store. One team is building a Payments service which has its own database; the service needs data that originates in the Accounts database. Both are using Amazon DynamoDB.

What approach will result in the simplest, decoupled, and reliable method to get near-real time updates from the Accounts database?

- A. Use Amazon Glue to perform frequent ETL updates from the Accounts database to the Payments database.
- B. Use Amazon ElastiCache in Payments, with the cache updated by triggers in the Accounts database.
- C. Use Amazon Kinesis Data Firehose to deliver all changes from the Accounts database to the Payments database.
- D. Use Amazon DynamoDB Streams to deliver all changes from the Accounts database to the Payments database.

Answer: D

Explanation:

Reference:

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

NEW QUESTION 76

- (Exam Topic 1)

You are writing to a DynamoDB table and receive the following exception: "ProvisionedThroughputExceededException". though according to your Cloudwatch metrics for the table, you are not exceeding your provisioned throughput.

What could be an explanation for this?

- A. You haven't provisioned enough DynamoDB storage instances
- B. You're exceeding your capacity on a particular Range Key
- C. You're exceeding your capacity on a particular Hash Key
- D. You're exceeding your capacity on a particular Sort Key
- E. You haven't configured DynamoDB Auto Scaling triggers

Answer: C

Explanation:

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

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

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-partition-key-design.html>

NEW QUESTION 80

- (Exam Topic 1)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 81

- (Exam Topic 1)

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

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

What is the SAFEST way to meet these requirements?

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

Answer: B

NEW QUESTION 85

- (Exam Topic 1)

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

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

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

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

Answer: AC

NEW QUESTION 90

- (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 93

- (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 94

- (Exam Topic 1)

A company maintains an application responsible for processing several thousand external callbacks each day. The company's System administrators want to know how many callbacks are being received on a rolling basis, and they want this data available for 10 days. The company also wants the ability to issue automated alerts if the number of callbacks exceeds the defined thresholds.

What is the MOST cost-effective way to address the need to track and alert on these statistics?

- A. Push callback data to an Amazon RDS database that can be queried to show historical data and to alert on exceeded thresholds.
- B. Push callback data to AWS X-Ray and use AWS Lambda to query, display, and alert on exceeded thresholds.
- C. Push callback data to Amazon Kinesis Data Streams and invoke an AWS Lambda function that stores data in Amazon DynamoDB and sends the required alerts.
- D. Push callback data to Amazon CloudWatch as a custom metric and use the CloudWatch alerting mechanisms to alert System Administrators.

Answer: D

NEW QUESTION 98

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 103

- (Exam Topic 1)

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

What is causing the duplicate entries?

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

Answer: B

Explanation:

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

NEW QUESTION 106

- (Exam Topic 1)

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

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

Answer: A

NEW QUESTION 107

- (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 109

- (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 110

- (Exam Topic 1)

An application running on EC2 instances is storing data in an S3 bucket. Security policy mandates that all data must be encrypted in transit. How can the Developer ensure that all traffic to the S3 bucket is encrypted?

- A. Install certificates on the EC2 instances.
- B. Create a bucket policy that allows traffic where SecureTransport is true.
- C. Create an HTTPS redirect on the EC2 instances.
- D. Create a bucket policy that denies traffic where SecureTransport is false.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-bucket-policies-and-apply-defense-in-depth-to-help-secure-y>

NEW QUESTION 113

- (Exam Topic 1)

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

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

Answer: AE

Explanation:

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

<https://docs.aws.amazon.com/AmazonS3/latest/dev/optimizing-performance.html>

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

NEW QUESTION 114

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 119

- (Exam Topic 1)

Which code snippet below returns the URL of a load balanced web site created in CloudFormation with an AWS::ElasticLoadBalancing::LoadBalancer resource name "ElasticLoad Balancer"?

- A. "Fn::Join" : ["http://", { "Fn::GetAtr" : ["ElasticLoadBalancer", "DNSName"] }]]]]
- B. "Fn::Join" : ["http://", { "Fn::GetAtr" : ["ElasticLoadBalancer", "Url"] }]]]]
- C. "Fn::Join" : ["http://", { "Ref" : "ElasticLoadBalancerUrl" }]]]]
- D. "Fn::Join" : ["http://", { "Ref" : "ElasticLoadBalancerDNSName" }]]]]

Answer: A

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-getatt.html>

NEW QUESTION 123

- (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 124

- (Exam Topic 1)

You have an environment that consists of a public subnet using Amazon VPC and 3 instances that are running in this subnet. These three instances can successfully communicate with other hosts on the Internet. You launch a fourth instance in the same subnet, using the same AMI and security group configuration you used for the others, but find that this instance cannot be accessed from the Internet.

What should you do to enable internet access?

- A. Deploy a NAT instance into the public subnet.
- B. Modify the routing table for the public subnet
- C. Configure a publically routable IP Address In the host OS of the fourth instance.
- D. Assign an Elastic IP address to the fourth instance.

Answer: C

NEW QUESTION 127

- (Exam Topic 1)

Developer is creating an AWS Lambda function to process a stream of data from an Amazon Kinesis Data Stream. When the Lambda function parses the data and encounters a missing field, it exits the function with an error. The function is generating duplicate records from the Kinesis stream. When the Developer looks at the stream output without the Lambda function, there are no duplicate records.

What is the reason for the duplicates?

- A. The Lambda function did not advance the Kinesis stream pointer to the next record after the error.
- B. The Lambda event source used asynchronous invocation, resulting in duplicate records.
- C. The Lambda function did not handle the error, and the Lambda service attempted to reprocess the data.
- D. The Lambda function is not keeping up with the amount of data coming from the stream.

Answer: C

Explanation:

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

NEW QUESTION 128

- (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 129

- (Exam Topic 2)

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

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

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

Answer: C

NEW QUESTION 131

- (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 134

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 138

- (Exam Topic 2)

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

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

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

Answer: A

NEW QUESTION 139

- (Exam Topic 2)

A developer is preparing a deployment package using AWS CloudFormation. 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 Ref 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 143

- (Exam Topic 2)

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

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

The Developer needs to create/delete branches.
Which specific IAM permissions need to be added, based on the principle of least privilege?

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

Answer: A

Explanation:

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

NEW QUESTION 148

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 149

- (Exam Topic 2)

An application writes items to an Amazon DynamoDB table. As the application scales to thousands of instances, calls to the DynamoDB API generate occasional ThrottlingException errors. The application is coded in a language incompatible with the AWS SDK.
How should the error be handled?

- A. Add exponential backoff to the application logic
- B. Use Amazon SQS as an API message bus
- C. Pass API calls through Amazon API Gateway
- D. Send the items to DynamoDB through Amazon Kinesis Data Firehose

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/throttled-ddb/>
SDKs automatically add exponential backoff. If not using the AWS SDKs, add your own backoff logic to the application code.
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Programming.Errors.html#Programming>.

NEW QUESTION 152

- (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 154

- (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 157

- (Exam Topic 2)

A company stores all personally identifiable information (PII) in an Amazon DynamoDB table named PII in Account A. An application running on Amazon EC2 instances in Account B requires access to the PII table. An administrator in Account A created an IAM role named AccessPII with privileges to access the PII table, and made account B a trusted entity.

Which combination of actional steps should Developers take to access the table? (Select TWO)

- A. Ask an Administrator in Account B to allow the EC2 IAM role permission to assume the AccessPII role
- B. Ask an Administrator in Account B to allow the EC2 IAM role permission to assume the AccessPII role with predefined service control policies
- C. Ask an Administrator in Account A to allow the EG2 IAM role permission to assume the AccessPII role with predefined service control policies.
- D. Include the AssumeRole API in the application code logic to obtain credentials to access the PII table
- E. Include the GetSession token API in the application code logic to obtain credentials to access the PII table

Answer: BE

NEW QUESTION 161

- (Exam Topic 2)

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

What should the Development team do to achieve this?

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

Answer: A

Explanation:

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

NEW QUESTION 163

- (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 164

- (Exam Topic 2)

A Developer is making changes to a custom application that is currently using AWS Elastic Beanstalk.

After the Developer completes the changes, what solutions will update the Elastic Beanstalk environment with the new application version? (Choose two.)

- A. Package the application code into a .zip file, and upload, then deploy the packaged application from the AWS Management Console
- B. Package the application code into a .tar file, create a new application version from the AWS Management Console, then update the environment by using AWS CLI
- C. Package the application code into a .tar file, and upload and deploy the packaged application from the AWS Management Console
- D. Package the application code into a .zip file, create a new application version from the packaged application by using AWS CLI, then update the environment by using AWS CLI
- E. Package the application code into a .zip file, create a new application version from the AWS Management Console, then rebuild the environment by using AWS CLI

Answer: CD

NEW QUESTION 165

- (Exam Topic 2)

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

Why is the table being throttled?

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

Answer: A

Explanation:

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

NEW QUESTION 170

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 171

- (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 174

- (Exam Topic 2)

An application ingests a large number of small messages and stores them in a database. The application uses AWS Lambda. A development team is making changes to the application's processing logic. In testing, it is taking more than 15 minutes to process each message. The team is concerned the current backend may time out.

Which changes should be made to the backend system to ensure each message is processed in the MOST scalable way?

- A. Add the messages to an Amazon SQS queue Set up an Amazon EC2 instance to poll the queue and process messages as they arrive.
- B. Add the messages to an Amazon SQS queue
- C. Set up Amazon EC2 instances in an Auto Scaling group to poll the queue and process the messages as they arrive.
- D. Create a support ticket to increase the Lambda timeout to 60 minutes to allow for increased processing time

E. Change the application to directly insert the body of the message into an Amazon RDS database.

Answer: A

NEW QUESTION 175

- (Exam Topic 2)

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

What is the fastest way to process all the files?

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

Answer: C

NEW QUESTION 177

- (Exam Topic 2)

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

How can the developer accomplish this?

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 181

- (Exam Topic 2)

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

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

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

Answer: A

NEW QUESTION 183

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 186

- (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 190

- (Exam Topic 2)

A developer implemented a static website hosted in Amazon S3 that makes web service requests in Amazon API Gateway and AWS Lambda. The site is showing an error that reads.

"No 'access control-allow-origin' header is present on the requested resource. Origin 'null is therefore not allowed access "
What should the developer do to resolve this issue?

- A. Enable cross-origin resource sharing (cors) on the s3 bucket
- B. Enable cross-origin resource sharing (cors) for the method in api gateway
- C. Add the access control-request-method header to the request
- D. Add the access-control-inquest headers header to the request

Answer: A

Explanation:

Reference: <https://forums.aws.amazon.com/thread.jspa?threadid=252972>

NEW QUESTION 193

- (Exam Topic 2)

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

Which AWS service should be used to accomplish this?

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

Answer: D

NEW QUESTION 195

- (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 200

- (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 205

- (Exam Topic 2)

The upload of a 15 GB object to Amazon S3 fails. The error message reads: "Your proposed upload exceeds the maximum allowed object size." What technique will allow the Developer to upload this object?

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

Answer: A

Explanation:

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

NEW QUESTION 208

- (Exam Topic 2)

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

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

F. Enable a DynamoDB stream, and trigger the Lambda function asynchronously from the stream.

Answer: C

Explanation:

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

NEW QUESTION 212

- (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 216

- (Exam Topic 2)

A developer is building an application that needs to store data in Amazon S3. Management requires that the data be encrypted before is it sent to Amazon S3 for storage. The encryption keys need to be managed by the security team.

Which approach should the developer take to meet these requirements?

- A. Implement server-side encryption using customer-provided encryption keys (SSE-C).
- B. Implement server-side encryption by using client-side master key.
- C. Implement client-side encryption using an AWS KMS managed customer master key (CMK).
- D. Implement Client-side encryption using Amazon S3 managed keys.

Answer: C

Explanation:

Reference: <https://aws.amazon.com/s3/faqs/>

NEW QUESTION 221

- (Exam Topic 2)

A Developer is investigating an issue whereby certain requests are passing through an Amazon API Gateway endpoint /MyAPI, but the requests do not reach the AWS Lambda function backing /MyAPI. The Developer found that a second Lambda function sometimes runs at maximum concurrency allowed for the given AWS account.

How can the Developer address this issue?

- A. Manually reduce the concurrent execution limit at the account level
- B. Add another API Gateway stage for /MyAPI, and shard the requests
- C. Configure the second Lambda function's concurrency execution limit
- D. Reduce the throttling limits in the API Gateway /MyAPI endpoint

Answer: C

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/11/set-concurrency-limits-on-individual-aws-lambda-functi> You can now set a concurrency limit on individual AWS Lambda functions. The concurrency limit you set will reserve a portion of your account level concurrency limit for a given function. This feature allows you to throttle a given function if it reaches a maximum number of concurrent executions allowed, which you can choose to set.

NEW QUESTION 224

- (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 229

- (Exam Topic 2)

A company is running a Docker application on Amazon ECS. The application must scale based on user load in the last 15 seconds.

How should a Developer instrument the code so that the requirement can be met?

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

Answer: B

Explanation:

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

NEW QUESTION 232

- (Exam Topic 2)

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

How should the developer accomplish this?

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

Answer: D

NEW QUESTION 236

- (Exam Topic 2)

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

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

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

Answer: A

NEW QUESTION 237

- (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 241

- (Exam Topic 2)

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

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

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

E. Count

Answer: BC

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-metrics-and-dimensions.html>

NEW QUESTION 243

- (Exam Topic 2)

A Development team would like to migrate their existing application code from a GitHub repository to AWS CodeCommit. What needs to be created before they can migrate a cloned repository to CodeCommit over HTTPS?

- A. A GitHub secure authentication token
- B. A public and private SSH key file
- C. A set of Git credentials generated from IAM
- D. An Amazon EC2 IAM role with CodeCommit permissions

Answer: C

Explanation:

<https://docs.aws.amazon.com/codecommit/latest/userguide/how-to-migrate-repository-existing.html>

NEW QUESTION 248

- (Exam Topic 2)

A developer must ensure that the IAM credentials used by an application in Amazon EC2 are not misused or compromised What should the developer use to keep user credentials secure?

- A. Environment variables
- B. AWS credentials file
- C. Instance profile credentials
- D. Command line options

Answer: C

NEW QUESTION 252

- (Exam Topic 3)

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

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

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

Answer: B

NEW QUESTION 257

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 261

- (Exam Topic 3)

A development team uses AWS Elastic Beanstalk for application deployment. The team has configured the application version lifecycle policy to limit the number of application versions to 25 However even with the lifecycle policy the source bundle is deleted from the Amazon S3 source bucket What should a developer do in the Elastic Beanstalk application version lifecycle settings to retain the source code in the S3 bucket?

- A. Change the Set the application versions limit by total count setting to zero.
- B. Disable the Lifecycle policy setting
- C. Change the Set the application version limit by age setting to zero.
- D. Set Retention to Retain source bundle in S3.

Answer: C

NEW QUESTION 264

- (Exam Topic 3)

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

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

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

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

Answer: C

NEW QUESTION 266

- (Exam Topic 3)

A developer is trying to get data from an Amazon DynamoDB table called demoman-table. The developer configured the AWS CLI to use a specific IAM user's credentials and executed the following command:

```
aws dynamodb get-item table-name demoman-table --key '{"id": <"N"; "1993"}' ' ' The command returned errors and no rows were returned
```

What is the MOST likely cause of these issues?

- A. The command is incorrect; it should be rewritten to use `aws dynamodb get-item --table demoman-table --key '{"id": "1993"}'` with a string argument.
- B. The developer needs to log a ticket with AWS Support to enable access to the demoman-table.
- C. Amazon DynamoDB cannot be accessed from the AWS CLI and needs to be called via the REST API.
- D. The IAM user needs an associated policy with read access to demoman-table.

Answer: D

NEW QUESTION 270

- (Exam Topic 3)

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

What is causing this issue?

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

Answer: B

NEW QUESTION 273

- (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 276

- (Exam Topic 3)

A physician's office management application requires that all data in transit between an EC2 instance and an Amazon EBS volume be encrypted. Which of the following techniques fulfills this requirement? (Select TWO)

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

Answer: AD

NEW QUESTION 277

- (Exam Topic 3)

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

- A. Deploy the WebSocket API to three stages for the respective routes: `create`, `update`, and `remove`.
- B. Create a new route key and set the name as `action`.

- C. Set the value of the route selection expression to action
- D. Set the value of the route selection expression to \$request.body.action

Answer: D

NEW QUESTION 279

- (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 284

- (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 286

- (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 287

- (Exam Topic 3)

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

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

Answer: AD

NEW QUESTION 292

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 293

- (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 295

- (Exam Topic 3)

What is required to trace Lambda-based applications with AWS X-Ray?

- A. Send logs from the Lambda application to an S3 bucket, trigger a Lambda function from that bucket to send data to AWS X-Ray.
- B. Trigger a Lambda function from the application logs in Amazon CloudWatch to submit tracing data to AWS X-Ray.
- C. Use an IAM execution role to give the Lambda function permissions and enable tracing.
- D. Update and add AWS X-ray daemon code to relevant parts of the Lambda function to set up the trace.

Answer: D

NEW QUESTION 299

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 304

- (Exam Topic 3)

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

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

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

Answer: D

NEW QUESTION 309

- (Exam Topic 3)

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

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

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

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

Answer: B

NEW QUESTION 310

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 314

- (Exam Topic 3)

An application uploads photos to an Amazon S3 bucket. Each photo that is uploaded to the S3 bucket must be resized to a thumbnail image by the application. Each thumbnail image is uploaded with a new name in the same S3 bucket.

Which AWS service can a developer configure to directly process each single S3 event (or each S3 object upload)?

- A. Amazon EC2
- B. Amazon Elastic Container Service (Amazon ECS)
- C. AWS Elastic Beanstalk
- D. AWS Lambda

Answer: D

NEW QUESTION 316

- (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 317

- (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 321

- (Exam Topic 3)

A developer is using Amazon DynamoDB to store application data. The developer wants to further improve application performance by reducing response times for read and write operations.

Which DynamoDB feature should be used to meet these requirements?

- A. Amazon DynamoDB Streams
- B. Amazon DynamoDB Accelerator
- C. Amazon DynamoDB global tables
- D. Amazon DynamoDB transactions

Answer: B

Explanation:

<https://aws.amazon.com/ko/blogs/database/amazon-dynamodb-accelerator-dax-a-read-throughwrite-through-cac>

NEW QUESTION 325

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 326

- (Exam Topic 3)

An application contains two components one component to handle HTTP requests, and another component to handle background processing tasks. Each component must scale independently. The developer wants to deploy this application using AWS Elastic Beanstalk. How should this application be deployed, based on these requirements?

- A. Deploy the application in a single Elastic Beanstalk environment
- B. Deploy each component in a separate Elastic Beanstalk environment
- C. Use multiple Elastic Beanstalk environments for the HTTP component but one environment for the background task component
- D. Use multiple Elastic Beanstalk environments for the background task component but one environment for the HTTP component

Answer: A

NEW QUESTION 330

- (Exam Topic 3)

A company is migrating the content delivery network for its dynamic PHP website to AWS. An Amazon CloudFront web distribution is part of the new infrastructure. The distribution has the following cache behavior settings:

- Allowed HTTP Methods is set to GET, HEAD
- Viewer Protocol Policy is set to HTTP and HTTPS

Developers test the solution and can reach the company's website over HTTP and HTTPS. However, the developers are unable to log in to the previously working administration panel of the website.

Which action will resolve this login issue?

- A. Set Allowed HTTP Methods to GET
- B. HEAD; OPTIONS
- C. Set Viewer Protocol Policy to HTTPS Only
- D. Set Allowed HTTP Methods to GET, HEAD, OPTIONS, PUT, POST, PATCH, DELETE
- E. Set Viewer Protocol Policy to Redirect HTTP to HTTPS

Answer: A

NEW QUESTION 335

- (Exam Topic 3)

A developer has code stored in an Amazon S3 bucket. The code must be deployed as an AWS Lambda function across multiple accounts in the same Region as the S3 bucket. The Lambda function will be deployed using an AWS CloudFormation template that is run for each account. What is the MOST secure approach to allow access to the Lambda code in the S3 bucket?

- A. Grant the CloudFormation execution role S3 list and get permissions. Add a bucket policy to Amazon S3 with the Principal of "AWS": [account numbers].
- B. Grant the CloudFormation execution role S3 get permissions. Add a bucket policy to Amazon S3 with the Principal of "".
- C. Use a service-based link to grant the Lambda function S3 list and get permissions by explicitly adding the S3 bucket's account number in the resource.
- D. Use a service-based link to grant the Lambda function S3 get permissions and add a Resource of "*" to allow access to the S3 bucket.

Answer: A

NEW QUESTION 336

- (Exam Topic 3)

A developer is leveraging a Border Gateway Protocol (BGP)-based AWS VPN connection to connect from on-premises to Amazon EC2 instances in the developer's account. The developer is able to access an EC2 instance in subnet A, but is unable to access an EC2 instance in subnet B in the same VPC. Which logs can the developer use to verify whether the traffic is reaching subnet B?

- A. VPN logs
- B. BGP logs
- C. VPC Flow Logs
- D. AWS CloudTrail logs

Answer: C

NEW QUESTION 338

- (Exam Topic 3)

A developer is building an AWS Lambda function that will dynamically generate and send a weekly newsletter to 100,000 users. This newsletter contains both static text and images. The developer needs a fast and highly scalable place to store the images that will be hyperlinked in the newsletter. Where should the developer store these images?

- A. Use an Amazon DynamoDB table with DynamoDB Streams and read capacity auto scaling enabled
- B. Use an Amazon S3 bucket and S3 Transfer Acceleration to speed up the image download
- C. Use an Amazon Aurora database with a public DNS endpoint and auto scaling enabled
- D. Use an Amazon S3 backed Amazon CloudFront distribution with a high Time-to-Live (TTL) to maximize caching

Answer: D

NEW QUESTION 339

- (Exam Topic 3)

An application running on multiple Amazon EC2 instances pulls messages ...SQS queue. A requirement for the application is that all messages must be encrypted at rest.

Developers are instructed to use methods that allow for centralized .. possible support requirements whenever possible.

Which of the following solution supports these requirements?

- A. Encrypt individual messages by using client-side encryption with customer managed keys, then write to the SQS queue.
- B. Encrypt individual messages by using SQS Extended Client and the Amazon S3 encryption client.
- C. Create an SQS queue, and encrypt the queue by using server-side encryption with AWS KMS
- D. Create an SQS queue and encrypt the queue by using client-side encryption

Answer: C

NEW QUESTION 343

- (Exam Topic 3)

A developer has built a market application that stores pricing data in Amazon DynamoDB with Amazon ElastiCache in front. The prices of items in the market change frequently. Sellers have begun complaining that, after they update the price of an item, the price does not actually change in the product listing. What could be causing this issue?

- A. The cache is not being invalidated when the price of the item is changed
- B. The price of the item is being retrieved using a write-through ElastiCache cluster
- C. The DynamoDB table was provisioned with insufficient read capacity
- D. The DynamoDB table was provisioned with insufficient write capacity.

Answer: A

NEW QUESTION 346

- (Exam Topic 3)

A developer is building a new application that uses an Amazon DynamoDB table. The specification states that all items that are older than 48 hours must be removed.

Which solution will meet this requirement?

- A. Create a new attribute that has the Number data type. Add a local secondary index (LSI) for this attribute and enable TTL with an expiration of 48 hours. In the application code, set the value of this attribute to the current timestamp for each new item that is being inserted.
- B. Create a new attribute that has the String data type. Add a local secondary index (LSI) for this attribute and enable TTL with an expiration of 48 hours. In the application code, set the value of this attribute to the current timestamp for each new item that is being inserted.
- C. Create a new attribute that has the Number data type. Enable TTL on the DynamoDB table for this attribute in the application code. Set the value of this attribute to the current timestamp plus 48 hours for each new item that is being inserted.
- D. Create a new attribute that has the String data type. Enable TTL on the DynamoDB table for this attribute. In the application code, set the value of this attribute to the current timestamp plus 48 hours for each new item that is being inserted.

Answer: C

NEW QUESTION 347

- (Exam Topic 3)

A photo sharing website gets millions of new images every week. The images are stored in Amazon S3 under a formatted date prefix. A developer wants to move images to a few S3 buckets for analysis and further processing. Images are not required to be moved in real time.

What is the MOST efficient method for performing this task?

- A. Use S3 PutObject events to invoke AWS Lambda. Then Lambda will copy the files to the other objects.
- B. Create an AWS Lambda function that will pull a day of images from the origin bucket and copy them to the other buckets.
- C. Use S3 Batch Operations to create jobs for images to be copied to each individual bucket.
- D. Use Amazon EC2 to batch pull images from multiple days and copy them to the other buckets.

Answer: D

NEW QUESTION 350

- (Exam Topic 3)

An IAM role is attached to an Amazon EC2 instance that explicitly denies access to all Amazon S3 API actions. The EC2 instance credentials file specifies the IAM access key and secret access key, which allow full administrative access.

Given that multiple modes of IAM access are present for this EC2 instance, which of the following is correct?

- A. The EC2 instance will only be able to list the S3 buckets.
- B. The EC2 instance will only be able to list the contents of one S3 bucket at a time.
- C. The EC2 instance will be able to perform all actions on any S3 bucket.
- D. The EC2 instance will not be able to perform any S3 action on any S3 bucket.

Answer: C

NEW QUESTION 351

- (Exam Topic 3)

A developer is building an application on Amazon EC2. The developer encountered an "Access Denied" error on some of the API calls to AWS services while testing. The developer needs to modify permissions that have been already given to the instance.

How can these requirements be met with minimal changes and minimum downtime?

- A. Make a new IAM role with the needed permissions. Stop the instance.
- B. Attach the new IAM role to the instance. Start the instance.

- C. Delete the existing 1AM role Attach a new 1AM role with the needed permissions
- D. Stop the instance Update the attached 1AM role adding the needed permission
- E. Start the instance
- F. Update the attached 1AM role adding the needed permissions

Answer: D

NEW QUESTION 352

- (Exam Topic 3)

A developer wants to secure sensitive configuration data such as passwords, database strings, and application license codes. Access to this sensitive information must be tracked for future audit purposes.

- A. In an encrypted file on the source code bundle; grant the application access with Amazon IAM
- B. In the Amazon EC2 Systems Manager Parameter Store; grant the application access with IAM
- C. On an Amazon EBS encrypted volume attach the volume to an Amazon EC2 instance to access the data
- D. As an object in an Amazon S3 bucket, grant an Amazon EC2 instance access with an IAM role.

Answer: B

NEW QUESTION 354

- (Exam Topic 3)

An application development team decides to use AWS X-Ray to monitor application code to analyze performance and perform root cause analysis. What does the team need to do to begin using X-Ray? (Select TWO)

- A. Log instrumentation output into an Amazon SQS queue
- B. Use a visualization tool to view application traces
- C. Instrument application code using the AWS SDK
- D. Install the X-Ray agent on the application servers
- E. Create an Amazon DynamoDB table to store the trace logs

Answer: DE

NEW QUESTION 356

- (Exam Topic 3)

A developer has written an AWS Lambda function using Java as the runtime environment. The developer wants to isolate a performance bottleneck in the code. Which steps should be taken to reveal the bottleneck?

- A. Use the Amazon CloudWatch API to write timestamps to a custom CloudWatch metric Use the CloudWatch console to analyze the resulting data
- B. Use the AWS X-Ray API to write trace data into X-Ray from strategic places within the code Use the Amazon CloudWatch console to analyze the resulting data
- C. Use the AWS X-Ray API to write trace data into X-Ray from strategic places within the code
- D. Use the X-Ray console to analyze the resulting data
- E. Use the Amazon CloudWatch API to write timestamps to a custom CloudWatch metric Use the AWS X-Ray console to analyze the resulting data

Answer: C

NEW QUESTION 358

- (Exam Topic 3)

A developer has written an application that runs on Amazon EC2 instances. The developer is adding functionality for the application to write objects to an Amazon S3 bucket. Which policy must the developer modify to allow the instances to write these objects?

- A. The IAM policy that is attached to the EC2 instance profile role
- B. The session policy that is applied to the EC2 instance role session
- C. The AWS Key Management Service (AWS KMS) key policy that is attached to the EC2 instance profile role.
- D. The Amazon VPC endpoint policy

Answer: A

NEW QUESTION 362

- (Exam Topic 3)

An application is experiencing performance issues based on increased demand. This increased demand is on read-only historical records pulled from an Amazon RDS-hosted database with custom views and queries. A developer improves performance without changing the database structure. Which approach will improve performance and MINIMIZE management overhead?

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

Answer: B

NEW QUESTION 364

- (Exam Topic 4)

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

- A. Launch Amazon EC2 instances for Windows Server

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

Answer: B

NEW QUESTION 368

- (Exam Topic 4)

A developer has an Amazon DynamoDB table that must be in provisioned mode to comply with user requirements. The application needs to support the following:

- Average item size: 10 KB
- Item reads each second: 10 strongly consistent
- Item writes each second: 2 transactional

Which read and write capacity cost-effectively meets these requirements?

- A. Read 10; write 2
- B. Read 30; write 40
- C. Use on-demand scaling
- D. Read 300; write 400

Answer: B

NEW QUESTION 373

- (Exam Topic 4)

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

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

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

Answer: A

NEW QUESTION 378

- (Exam Topic 4)

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

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

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

Answer: B

NEW QUESTION 380

- (Exam Topic 4)

A company has a new application. The company needs to secure sensitive configuration data such as database connection strings, application license codes, and API keys that the application uses to access external resources. The company must track access to the configuration data for auditing purposes. The resources are managed outside the application.

The company is not required to manage rotation of the connection strings, license codes, and API keys in the application. The company must implement a solution to securely store the configuration data and to give the application access to the configuration data. The solution must comply with security best practices.

Which solution will meet these requirements MOST cost-effectively?

- A. Store the configuration data in an encrypted file on the source code bundl
- B. Grant the application access by using IAM policies.
- C. Store the configuration data in AWS Systems Manager Parameter Stor
- D. Grant the application access by using IAM policies.
- E. Store the configuration data on an Amazon Elastic Block Store (Amazon EBS) encrypted volume. Attach the EBS volume to an Amazon EC2 instance to provide the application with access to the data.
- F. Store the configuration data in AWS Secrets Manage

G. Grant the application access by using IAM policies.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/mt/the-right-way-to-store-secrets-using-parameter-store/> <https://docs.aws.amazon.com/managedservices/latest/userguide/sys-man-param-store.html> AWS Systems Manager Parameter Store (AMS SSPS):

AWS Systems Manager Parameter Store provides secure, hierarchical storage for configuration data management and secrets management. You can store data such as passwords, database strings, and license codes as parameter values.

NEW QUESTION 381

- (Exam Topic 4)

A company has an application that provides blog hosting services to its customers. The application includes an Amazon DynamoDB table with a primary key. The primary key consists of the customers' Username as a partition key and the NumberOfBlogs as a sort key. The application stores the TotalReactionsOnBlogs as an attribute on the same DynamoDB table.

A developer needs to implement an operation to retrieve the top 10 customers based on the greatest number of reactions on their blogs. This operation must not consume the DynamoDB table's existing read capacity.

What should the developer do to meet these requirements in the MOST operationally efficient manner?

- A. For the existing DynamoDB table, create a new global secondary index (GSI) that has the Username as a partition key and the TotalReactionsOnBlogs as a sort key.
- B. For the existing DynamoDB table, create a new local secondary index (LSI) that has the Username as a partition key and the TotalReactionsOnBlogs as a sort key.
- C. Back up and restore the DynamoDB table to a new DynamoDB table. Create a new global secondary index (GSI) that has the Username as a partition key and the TotalReactionsOnBlogs as a sort key. Delete the old DynamoDB table.
- D. Back up and restore the DynamoDB table to a new DynamoDB table.
- E. Create a new local secondary index (LSI) that has the Username as a partition key and the TotalReactionsOnBlogs as a sort key.
- F. Delete the old DynamoDB table.

Answer: B

NEW QUESTION 384

- (Exam Topic 4)

A developer is troubleshooting a new AWS Lambda function. The function should run automatically each time a new object is uploaded to an Amazon S3 bucket. However, the developer finds that all calls failed before they reached the application code inside the function.

Which of the following is a possible reason for this failure?

- A. The function resource policy does not allow access from Amazon S3.
- B. The function execution role does not allow access from Amazon S3.
- C. The function execution role does not allow access to Amazon S3.
- D. The IAM user does not have access to Amazon S3.

Answer: C

NEW QUESTION 389

- (Exam Topic 4)

A developer is exposing an API by using Amazon API Gateway and AWS Lambda as the backend for an application. The developer wants to add validation rules for a POST method to ensure that the data (from the frontend web form) is valid. The validation rules must include mandatory fields, data type, length, and regular expressions.

Which solution will meet these requirements?

- A. Create an API Gateway model with schema for data validation.
- B. Create API Gateway HTTP request headers for data validation.
- C. Create API Gateway URL query string parameters for data validation.
- D. Create API Gateway URL path parameters for data validation.

Answer: D

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-method-request-validation.html>

NEW QUESTION 392

- (Exam Topic 4)

A developer is testing a new file storage application that uses an Amazon CloudFront distribution to serve content from an Amazon S3 bucket. The distribution accesses the S3 bucket by using an origin access identity (OAI). The S3 bucket's permissions explicitly deny access to all other users.

The application prompts users to authenticate on a login page and then uses signed cookies to allow users to access their personal storage directories. The developer has configured the distribution to use its default cache behavior with restricted viewer access and has set the origin to point to the S3 bucket. However, when the developer tries to navigate to the login page, the developer receives a 403 Forbidden error.

The developer needs to implement a solution to allow unauthenticated access to the login page. The solution also must keep all private content secure.

Which solution will meet these requirements?

- A. Add a second cache behavior to the distribution with the same origin as the default cache behavior.
- B. Set the path pattern for the second cache behavior to the path of the login page, and make viewer access unrestricted.
- C. Keep the default cache behavior's settings unchanged.
- D. Add a second cache behavior to the distribution with the same origin as the default cache behavior.
- E. Set the path pattern for the second cache behavior to *, and make viewer access restricted.
- F. Change the default cache behavior's path pattern to the path of the login page, and make viewer access unrestricted.
- G. Add a second origin as a failover origin to the default cache behavior.
- H. Point the failover origin to the S3 bucket.

- I. Set the path pattern for the primary origin to * and make viewer access restricted
- J. Set the path pattern for the failover origin to the path of the login page, and make viewer access unrestricted.
- K. Add a bucket policy to the S3 bucket to allow read access
- L. Set the resource on the policy to the Amazon Resource Name (ARN) of the login page object in the S3 bucket
- M. Add a CloudFront function to the default cache behavior to redirect unauthorized requests to the login page's S3 URI.

Answer: B

Explanation:

Adding a second cache behavior to the distribution with the same origin as the default cache behavior and setting the path pattern to * will allow access to all files in the S3 bucket. Changing the default cache behavior's path pattern to the path of the login page and making viewer access unrestricted will allow unauthenticated users to access the login page, while keeping all other private content secure.

NEW QUESTION 395

- (Exam Topic 4)

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

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

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 400

- (Exam Topic 4)

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

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

Answer: A

NEW QUESTION 401

- (Exam Topic 4)

A development team set up a pipeline to launch a test environment. The developers want to automate tests for their application. The team created an AWS CodePipeline stage to deploy the application to a test environment in batches using AWS Elastic Beanstalk. A later CodePipeline stage contains a single action that uses AWS CodeBuild to run numerous automated Selenium-based tests on the deployed application. The team must speed up the pipeline without removing any of the individual tests.

Which set of actions will MOST effectively speed up application deployment and testing?

- A. Set up an all-at-once deployment in Elastic Beanstalk
- B. Run tests in parallel with multiple CodeBuild actions.
- C. Set up a rolling update in Elastic Beanstalk
- D. Run tests in serial with a single CodeBuild action.
- E. Set up an immutable update in Elastic Beanstalk
- F. Run tests in serial with a single CodeBuild action.
- G. Set up a traffic-splitting deployment in Elastic Beanstalk
- H. Run tests in parallel with multiple CodeBuild actions.

Answer: A

Explanation:

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

All at once – The quickest deployment method. Suitable if you can accept a short loss of service, and if quick deployments are important to you. With this method, Elastic Beanstalk deploys the new application version to each instance. Then, the web proxy or application server might need to restart. As a result, your application might be unavailable to users (or have low availability) for a short time.

NEW QUESTION 403

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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