



Amazon

Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate

About ExamBible

Your Partner of IT Exam

Found in 1998

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

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

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

NEW QUESTION 1

- (Exam Topic 1)

A Developer 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 2

- (Exam Topic 1)

For a deployment using AWS CodeDeploy, what is the run order of the hooks for in-place deployments?

- A. Before Install -> Application Stop -> Application Start -> After Install
- B. Application Stop -> Before Install -> After Install -> Application Start
- C. Before Install -> Application Stop -> Validate Service -> Application Start
- D. Application Stop -> Before Install -> Validate Service -> Application Start

Answer: B

NEW QUESTION 3

- (Exam Topic 1)

You attempt to store an object in the US-STANDARD region in Amazon S3, and receive a confirmation that it has been successfully stored. You then immediately make another API call and attempt to read this object. S3 tells you that the object does not exist. What could explain this behavior?

- A. US-STANDARD uses eventual consistency and it can take time for an object to be readable in a bucket
- B. Objects in Amazon S3 do not become visible until they are replicated to a second region.
- C. US-STANDARD imposes a 1 second delay before new objects are readable.
- D. You exceeded the bucket object limit, and once this limit is raised the object will be visible.

Answer: A

Explanation:

<https://acloud.guru/forums/aws-certified-developer-associate/discussion/-KGngHzVQ03OpeAA9jSP/i-cant-ans> https://acloud.guru/forums/aws-certified-developer-associate/discussion/-K5WKXRAIJdOu58GREF_/s3-questio

NEW QUESTION 4

- (Exam Topic 1)

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

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

Answer: D

NEW QUESTION 5

- (Exam Topic 1)

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

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

Answer: BC

Explanation:

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

NEW QUESTION 6

- (Exam Topic 1)

A startup's photo-sharing site is deployed in a VPC. An ELB distributes web traffic across two subnets. ELB session stickiness is configured to use the AWS-generated session cookie, with a session TTL of 5 minutes. The webserver Auto Scaling Group is configured as: min-size=4, max-size=4. The startups preparing for a public launch, by running load-testing software installed on a single EC2 instance running in us-west-2a. After 60 minutes of load-testing, the webserver logs show:

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

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

Answer: CE

NEW QUESTION 7

- (Exam Topic 1)

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

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

Answer: E

Explanation:

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

NEW QUESTION 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)

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

How can these requirements be met?

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

- J. Store the source code in AWS CodeComm
- K. Create a CodePipeline to automate unit testin
- L. Use Amazon CloudWatch to trigger notification of failure events.

Answer: D

NEW QUESTION 10

- (Exam Topic 1)

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

- A. Define a Swagger fil
- B. Use AWS Elastic Beanstalk to deploy the Swagger file.
- C. Define a Swagger fil
- D. Use AWS CodeDeploy to deploy the Swagger file.
- E. Deploy a SAM template with an inline Swagger definition.
- F. Define a Swagger fil
- G. Deploy a SAM template that references the Swagger file.
- H. Define an inline Swagger definition in a Lambda functio
- I. Invoke the Lambda function.

Answer: CD

Explanation:

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

NEW QUESTION 12

- (Exam Topic 1)

An AWS Elastic Beanstalk application needs to be deployed in multiple regions and requires a different Amazon Machine Image (AMI) in each region. Which AWS CloudFormation template key can be used to specify the correct AMI for each region?

- A. Parameters
- B. Outputs
- C. Mappings
- D. Resources

Answer: C

Explanation:

Reference: <https://docs.aws.amazon.com/marketplace/latest/userguide/cloudformation.html>

NEW QUESTION 16

- (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 19

- (Exam Topic 1)

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

What services can the Developer use to meet these requirements?

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

Answer: B

Explanation:

AWS documentation - Cognito MFA Managing Security

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

NEW QUESTION 22

- (Exam Topic 1)

A company uses Amazon DynamoDB for managing and tracking orders. The DynamoDB table is partitioned based on the order date. The company receives a huge increase in orders during a sales event, causing DynamoDB writes to throttle, and the consumed throughput is far below the provisioned throughput. According to AWS best practices, how can this issue be resolved with MINIMAL costs?

- A. Create a new DynamoDB table for every order date.
- B. Increase the read and write capacity units of the DynamoDB table.
- C. Add a random number suffix to the partition key values.
- D. Add a global secondary index to the DynamoDB table.

Answer: C

Explanation:

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

NEW QUESTION 27

- (Exam Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 32

- (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 36

- (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 41

- (Exam Topic 1)

An application is real-time processing millions of events that are received through an API. What service could be used to allow multiple consumers to process the data concurrently and MOST cost-effectively?

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

D. Amazon Kinesis Streams

Answer: D

NEW QUESTION 45

- (Exam Topic 1)

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

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

Answer: B

NEW QUESTION 50

- (Exam Topic 1)

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

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

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

Answer: BE

NEW QUESTION 54

- (Exam Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 55

- (Exam Topic 1)

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

What can the Developer do to fix this problem?

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

Answer: C

NEW QUESTION 59

- (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 64

- (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 69

- (Exam Topic 1)

Which of the following statements about SQS is true?

- A. Messages will be delivered exactly once and messages will be delivered in First in, First out order
- B. Messages will be delivered exactly once and message delivery order is indeterminate
- C. Messages will be delivered one or more times and messages will be delivered in First in, First out order
- D. Messages will be delivered one or more times and message delivery order is indeterminate

Answer: D

Explanation:

<https://aws.amazon.com/sqs/features/>

NEW QUESTION 72

- (Exam Topic 1)

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

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

Answer: C

Explanation:

Refer AWS documentation - SAM Gradual Code Deployment

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

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

NEW QUESTION 76

- (Exam Topic 1)

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

How should the application authenticate with AWS services in production?

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

Answer: A

Explanation:

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

NEW QUESTION 81

- (Exam Topic 1)

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

You advise them to:

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

Answer: D

Explanation:

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

NEW QUESTION 86

- (Exam Topic 1)

What are the steps to using the AWS CLI to launch a templated serverless application?

- A. Use AWS CloudFormation get-template then CloudFormation execute-change-set.
- B. Use AWS CloudFormation validate-template then CloudFormation create-change-set.
- C. Use AWS CloudFormation package then CloudFormation deploy.
- D. Use AWS CloudFormation create-stack then CloudFormation update-stack.

Answer: C

Explanation:

<https://docs.aws.amazon.com/cli/latest/reference/cloudformation/package.html>

NEW QUESTION 90

- (Exam Topic 1)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 94

- (Exam Topic 1)

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

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

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

Answer: C

NEW QUESTION 99

- (Exam Topic 1)

A company is creating an application that will require users to access AWS services and allow them to reset their own passwords.

Which of the following would allow the company to manage users and authorization while allowing users to reset their own passwords?

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

Answer: D

Explanation:

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

NEW QUESTION 103

- (Exam Topic 1)

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

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

Answer: BD

NEW QUESTION 107

- (Exam Topic 1)

A company has multiple Developers located across the globe who are updating code incrementally for a development project. When Developers upload code concurrently, internet connectivity is slow, and it is taking a long time to upload code for deployment in AWS Elastic Beanstalk. Which step will result in minimized upload and deployment time with the LEAST amount of administrative effort?

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

Answer: B

Explanation:

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

NEW QUESTION 111

- (Exam Topic 1)

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

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

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

- A. Convert the file to JSON format.
- B. Rename the file to a .config extension.
- C. Change the configuration section from options_settings to resources.
- D. Change the namespace of the option settings to a custom namespace.

Answer: B

Explanation:

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

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

NEW QUESTION 114

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 117

- (Exam Topic 1)

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

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

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

Answer: C

NEW QUESTION 121

- (Exam Topic 1)

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

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

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

Answer: D

NEW QUESTION 123

- (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 124

- (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 126

- (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 127

- (Exam Topic 1)

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

A client error (UnauthorizedOperation) occurred when calling the RunInstances operation: You are not authorized to perform this operation. Encoded authorization failure message: oGsbAaIV7wlfj8zUqebHUANHzFbmkzILlxyj_y9xwhIHK99U_cUq1FIeZnskWDjQ1wSHStVfdCEyZILGocccGpCiC IhORceWF9rRwFTnEcRJ3N9iTrPAE1WHveC5Z54ALPaWlEjHlLg8CaB8d81CKmxQuylCm0r1Bf2fHJRujAYopMvmga 8o1FmKAl9yn_Z5rI120Q9p5ZIMX28zYM4dTulcJQUQjosgrEejfiIMYDda8170oko9H6VmGJX62KfkRa517yE6hhh 2bIwA6tpyCJy2LWFRTe4bafqAyoqkarhPA4mGiZyWn4gSqbo8o- uqSivWYPweaKGkampa0arcFR4gBD7Ph097WYBkzX9hVjGppLMy4jpXRv

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 131

- (Exam Topic 1)

A company has written a Java AWS Lambda function to be triggered whenever a user uploads an image to an Amazon S3 bucket. The function converts the original image to several different formats and then copies the resulting images to another Amazon S3 bucket. The Developers find that no images are being copied to the second Amazon S3 bucket. They have tested the code on an Amazon EC2 instance with 1GB of RAM, and it takes an average of 500 seconds to complete. What is the MOST likely cause of the problem?

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

Answer: C

NEW QUESTION 134

- (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 137

- (Exam Topic 1)

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

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

Answer: D

NEW QUESTION 142

- (Exam Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 146

- (Exam Topic 1)

A serverless application uses an API Gateway and AWS Lambda.

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

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

Answer: A

NEW QUESTION 151

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 153

- (Exam Topic 1)

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

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

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

Answer: A

NEW QUESTION 156

- (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 157

- (Exam Topic 1)

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

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

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

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

```
    db.close()
```

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

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

Answer: C

Explanation:

Refer AWS documentation - Lambda Best Practices

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

NEW QUESTION 161

- (Exam Topic 1)

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

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

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

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

Answer: DE

NEW QUESTION 166

- (Exam Topic 1)

A large e-commerce site is being designed to deliver static objects from Amazon S3. The Amazon S3 bucket will serve 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 171

- (Exam Topic 1)

In DynamoDB, what type of HTTP response codes indicate that a problem was found with the client request sent to the service?

- A. 5xx HTTP response code
- B. 200 HTTP response code
- C. 306 HTTP response code
- D. 4xx HTTP response code

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/API/ErrorResponses.html#ErrorCodeList>

NEW QUESTION 174

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 177

- (Exam Topic 1)

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

What is the recommended solution?

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

Answer: D

NEW QUESTION 181

- (Exam Topic 1)

How is provisioned throughput affected by the chosen consistency model when reading data from a DynamoDB table?

- A. Strongly consistent reads use the same amount of throughput as eventually consistent reads
- B. Strongly consistent reads use more throughput than eventually consistent reads.
- C. Strongly consistent reads use less throughput than eventually consistent reads
- D. Strongly consistent reads use variable throughput depending on read activity

Answer: B

NEW QUESTION 182

- (Exam Topic 1)

A Developer is designing a new application that uses Amazon S3. To satisfy compliance requirements, the Developer must encrypt the data at rest. How can the Developer accomplish this?

- A. Use s3:x-amz-acl as a condition in the S3 bucket policy.
- B. Use Amazon RDS with default encryption.
- C. Use aws:SecureTransport as a condition in the S3 bucket policy.
- D. Turn on S3 default encryption for the S3 bucket.

Answer: D

NEW QUESTION 186

- (Exam Topic 2)

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

Which approach would provide the BEST user experience?

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

Answer: C

NEW QUESTION 187

- (Exam Topic 2)

A company is managing a NoSQL database on-premises to host a critical component of an application, which is starting to have scaling issues. The company wants to migrate the application to Amazon DynamoDB with the following considerations:

- Optimize frequent queries
- Reduce read latencies
- Plan for frequent queries on certain key attributes of the table Which solution would help achieve these objectives?

- A. Create global secondary indexes on keys that are frequently queried Add the necessary attributes into the indexes.
- B. Create local secondary indexes on keys that are frequently queried DynamoDB will fetch needed attributes from the table .
- C. Create DynamoDB global tables to speed up query responses Use a scan to fetch data from the table.
- D. Create an AWS Auto Scaling policy for the DynamoDB table

Answer: A

Explanation:

"Global secondary index—An index with a partition key and a sort key that can be different from those on the base table. A global secondary index is considered "global" because queries on the index can span all of the data in the base table, across all partitions.

Local secondary index—An index that has the same partition key as the base table, but a different sort key. A local secondary index is "local" in the sense that every partition of a local secondary index is scoped to a base table partition that has the same partition key value. "

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

NEW QUESTION 188

- (Exam Topic 2)

A Developer created a new AWS account and must create a scalable AWS Lambda function that meets the following requirements for concurrent execution:

- > Average execution time of 100 seconds
- > 50 requests per second

Which step must be taken prior to deployment to prevent errors?

- A. Implement dead-letter queues to capture invocation errors
- B. Add an event source from Amazon API Gateway to the Lambda function
- C. Implement error handling within the application code
- D. Contact AWS Support to increase the concurrent execution limits

Answer: D

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/05/aws-lambda-raises-default-concurrent-execution-limit/>

NEW QUESTION 189

- (Exam Topic 2)

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

What is the best option to store the session state?

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

Answer: A

Explanation:

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

NEW QUESTION 192

- (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 196

- (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 200

- (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 201

- (Exam Topic 2)

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

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

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

Answer: D

NEW QUESTION 205

- (Exam Topic 2)

A Developer is trying to deploy a serverless application using AWS CodeDeploy. The application was updated and needs to be redeployed. What file does the Developer need to update to push that change through CodeDeploy?

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

Answer: C

Explanation:

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

NEW QUESTION 207

- (Exam Topic 2)

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

Which caching strategy will address this situation MOST efficiently?

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

Answer: B

Explanation:

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

NEW QUESTION 211

- (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-provisioning-write-capacity-for-global-secondary-index> <https://medium.com/@synchrophoto/amazon-dynamodb-provisioning-write-capacity-for-global-secondary-index>

NEW QUESTION 216

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 220

- (Exam Topic 2)

A developer has written an Amazon Kinesis Data Streams application. As usage grows and traffic increases over time, the application is regularly receiving ProvisionedThroughputExceededException error messages. Which steps should the developer take to resolve the error? (Select TWO.)

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

Answer: BC

Explanation:

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

NEW QUESTION 223

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 225

- (Exam Topic 2)

An application runs on multiple EC2 instances behind an ELB. Where is the session data best written so that it can be served reliably across multiple requests?

- A. Write data to Amazon ElastiCache
- B. Write data to Amazon Elastic Block Store.
- C. Write data to Amazon EC2 Instance Store.
- D. Write data to the root filesystem.

Answer: C

Explanation:

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

NEW QUESTION 228

- (Exam Topic 2)

A developer is writing an application that will process data delivered into an Amazon S3 bucket. The data is delivered approximately 10 times a day, and the developer expects the data will be processed in less than 1 minute, on average. How can the developer deploy and invoke the application with the lowest cost and lowest latency?

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

Answer: A

Explanation:

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

NEW QUESTION 233

- (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 237

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 241

- (Exam Topic 2)

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

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

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

Answer: B

NEW QUESTION 243

- (Exam Topic 2)

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

Which approach will meet these requirements?

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

Answer: D

Explanation:

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

NEW QUESTION 248

- (Exam Topic 2)

A company caches session information for a web application in an Amazon DynamoDB table. The company wants an automated way to delete old items from the table.

What is the simplest way to do this?

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

Answer: B

Explanation:

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

NEW QUESTION 252

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 255

- (Exam Topic 2)

A developer is updating an application deployed on AWS Elastic Beanstalk. The new version is incompatible with the old version. To successfully deploy the update, a full cutover to the new updated version must be performed on all instances at one time, with the ability to roll back changes in case of a deployment failure in the new version.

How can this be performed with the LEAST amount of downtime?

- A. Use the Elastic Beanstalk All at once deployment policy to update all instances simultaneously.
- B. Perform an Elastic Beanstalk Rolling with additional batch deployment.
- C. Deploy the new version in a new Elastic Beanstalk environment and swap environment URLs.
- D. Perform an Elastic Beanstalk Rolling deployment.

Answer: C

NEW QUESTION 259

- (Exam Topic 2)

An application running on an Amazon Linux EC2 instance needs to manage the AWS infrastructure. How can the EC2 instance be configured to make AWS API calls securely?

- A. Sign the AWS CLI command using the signature version 4 process.
- B. Run the `aws configure` AWS CLI command and specify the access key id and secret access key.
- C. Specify a role for the EC2 instance with the necessary privileges.
- D. Pass the access key id and secret access key as parameters for each AWS CLI command.

Answer: C

NEW QUESTION 263

- (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 268

- (Exam Topic 2)

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

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

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

Answer: AB

Explanation:

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

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

NEW QUESTION 271

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 273

- (Exam Topic 2)

A developer is setting up Amazon API gateway for their company's products. The API will be registered developers to query and update their environments. The company wants to limit the amount of requests end users send for.bot cost and security reason management wants to offer registered the option of buying larger packages that allow for more requests.

- A. Enable throttling for the API Gateway stage Set a value for both the rate and burst capacity If a registered larger package, create a stage for them, adjust the values, and share the new URL with them.
- B. 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
- C. 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
- D. 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: A

NEW QUESTION 275

- (Exam Topic 2)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 276

- (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 280

- (Exam Topic 2)

A company needs to distribute firmware updates to its customers around the world. Which service will allow easy and secure control of the access to the downloads at the lowest cost?

- A. Use Amazon CloudFront with signed URLs for Amazon S3
- B. Create a dedicated Amazon CloudFront Distribution for each customer
- C. Use Amazon CloudFront with AWS Lambda@Edge
- D. Use Amazon API Gateway and AWS Lambda to control access to an S3 bucket

Answer: A

Explanation:

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

NEW QUESTION 283

- (Exam Topic 2)

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

Which solution will meet this requirement?

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

Answer: C

NEW QUESTION 288

- (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 291

- (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 it is 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 293

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 295

- (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 298

- (Exam Topic 2)

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

How can DynamoDB costs be minimized while maximizing application performance?

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

Answer: C

Explanation:

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

NEW QUESTION 303

- (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 additional 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 EC2 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 GetSessionToken API in the application code logic to obtain credentials to access the PII table

Answer: AD

NEW QUESTION 308

- (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 313

- (Exam Topic 2)

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

How can the Developer meet this requirement?

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

Answer: C

NEW QUESTION 318

- (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 319

- (Exam Topic 2)

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

Which approach should the Developer use to achieve this goal?

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

Answer: C

NEW QUESTION 322

- (Exam Topic 2)

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

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

- A. Run the following: `docker pull REPOSITORY_URI : TAG`
- B. Run the output of the following: `aws ecr get-login` and then run: `docker pull REPOSITORY_URI : TAG`
- C. Run the following: `aws ecr get-login` and then run: `docker pull REPOSITORY_URI : TAG`
- D. Run the output of the following: `aws ecr get-download-url-for-layer` and then run: `docker pull REPOSITORY_URI : TAG`

Answer: B

Explanation:

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

NEW QUESTION 325

- (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 326

- (Exam Topic 2)

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

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

- A. Increase the backend T2 EC2 instance size to x1 to handle the largest possible load throughout the year
- B. Implement an Amazon SQS queue to decouple the front-end and backend servers
- C. Use an Amazon SNS queue to decouple the front-end and backend servers.
- D. Migrate the backend servers to on-premises and pull from an Amazon SNS queue
- E. Modify the backend servers to pull from an Amazon SQS queue.

Answer: CD

NEW QUESTION 327

- (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 328

- (Exam Topic 3)

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

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

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

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

Answer: D

Explanation:

Meaning:

DOC-EXAMPLE-BUCKET ==> bucket

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

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

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

NEW QUESTION 332

- (Exam Topic 3)

A developer has written an application that uses Amazon API Gateway and AWS Lambda The developer needs to configure the application so that the developer can visualize the application's components and Identify performance bottlenecks

What should the developer do to meet these requirements?

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

Answer: A

NEW QUESTION 336

- (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 339

- (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 343

- (Exam Topic 3)

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

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

Answer: BE

NEW QUESTION 345

- (Exam Topic 3)

A developer needs to manage AWS infrastructure as code and must be able to deploy multiple identical copies of the infrastructure, stage changes, and revert to previous versions.

Which approach addresses these requirements?

- A. Use cost allocation reports and AWS OpsWorks to deploy and manage the infrastructure.
- B. Use Amazon CloudWatch metrics and alerts along with resource tagging to deploy and manage the infrastructure.
- C. Use AWS Elastic Beanstalk and AWS CodeCommit to deploy and manage the infrastructure.
- D. Use AWS CloudFormation and AWS CodeCommit to deploy and manage the infrastructure.

Answer: D

NEW QUESTION 347

- (Exam Topic 3)

A developer is writing a new AWS Serverless Application Model (AWS SAM) template with a new AWS Lambda function. The Lambda function runs complex code. The developer wants to test the Lambda function with more CPU power.

What should the developer do to meet this requirement?

- A. Increase the runtime engine version
- B. Increase the timeout
- C. Increase the number of Lambda layers.
- D. Increase the memory

Answer: D

NEW QUESTION 349

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 350

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 353

- (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 355

- (Exam Topic 3)

A developer wants to insert a record into an Amazon DynamoDB table as soon as a new file is added to an Amazon S3 bucket. Which set of steps would be necessary to achieve this?

- A. Create an event with Amazon CloudWatch Events that will monitor the S3 bucket and then insert the records into DynamoDB
- B. Configure an S3 event to invoke a Lambda function that inserts records into DynamoDB
- C. Create a Lambda function that will poll the S3 bucket and then insert the records into DynamoDB.
- D. Create a cron job that will run at a scheduled time and insert the records into DynamoDB

Answer: B

NEW QUESTION 356

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 360

- (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 364

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 366

- (Exam Topic 3)

A development team uses AWS Elastic Beanstalk to deploy a Java-based web application. The team wants to ensure that changes to the source code and the configuration are always deployed on new instances. The team configures the Elastic Beanstalk environment to use immutable updates. However, an error occurs the first time a change is deployed with the new update policy. What is the MOST likely cause of this issue?

- A. Immutable updates are not supported for Java-based applications.
- B. The account has reached its on-demand instance limit.
- C. Immutable updates are only supported for m4 large and larger instance types.
- D. The developer must also modify the `ebextensions/immutable-updates` config file to enable immutable updates.

Answer: A

NEW QUESTION 368

- (Exam Topic 3)

A developer is using AWS CodeDeploy to deploy an application running on Amazon EC2. The developer wants to change the file permissions for a specific deployment file. Which lifecycle event should a developer use to meet this requirement?

- A. AfterInstall
- B. DownloadBundle

- C. BeforeInstall
- D. ValdateService

Answer: A

NEW QUESTION 371

- (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 376

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 377

- (Exam Topic 3)

A developer is working on an e-commerce 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 382

- (Exam Topic 3)

A developer is developing an application that uses signed requests (Signature Version 4) to call other AWS services. The developer has created a canonical request, has created the string to sign, and has calculated signing information. Which methods could the developer use to complete a signed request? (Select TWO)

- A. Add the signature to an HTTP header that is named Authorization
- B. Add the signature to a session cookie
- C. Add the signature to an HTTP header that is named Authentication
- D. Add the signature to a query string parameter that is named X-Amz-Signature
- E. Add the signature to an HTTP header that is named WWW-Authenticate

Answer: DE

NEW QUESTION 383

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 384

- (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 385

- (Exam Topic 3)

A development team is building a new application that will run on Amazon EC2 and use Amazon DynamoDB as a storage layer The developers all have assigned IAM user accounts in the same IAM group The developers currently can launch EC2 instances and they need to be able to launch EC2 instances with an instance role allowing access to Amazon DynamoDB.

Which AWS IAM changes are needed when creating an instance role to provide this functionality?

- A. Create an IAM permission policy attached to the role that allows access to DynamoDB Add a trust policy to the role that allows DynamoDB to assume the role Attach a permissions policy to the development group in AWS IAM that allows developers to use the IAM GetRole and IAM PassRole permissions for the role
- B. Create an IAM permissions policy attached to the role that allows access to DynamoDB Add a trust policy to the role that allows Amazon EC2 to assume the role Attach a permissions policy to the development group in AWS IAM that allows developers to use the IAM PassRole permission for the role
- C. Create an IAM permission policy attached to the role that allows access to Amazon EC2 Add a trust policy to the role that allows DynamoDB to assume the role Attach a permissions policy to the development group in AWS IAM that allows developers to use the IAM PassRole permission for the role
- D. Create an IAM permissions policy attached to the role that allows access to DynamoDB Add a trust policy to the role that allows Amazon EC2 to assume the role Attach a permissions policy to the development group in AWS IAM that allows developers to use the iam GetRole permission for the role.

Answer: C

NEW QUESTION 387

- (Exam Topic 3)

An application is processing clickstream data using Amazon Kinesis. The clickstream data feed into Kinesis experiences periodic spikes. The PutRecords API call occasionally fails and the logs show that the failed call returns the response shown below.

```
{
  "FailedRecordCount": 1,
  "Records": [
    {
      "SequenceNumber": "212693199899900637946712965403778482371",
      "ShardId": "shardId-000000000001"
    },
    {
      "ErrorCode": "ProvisionedThroughputExceededException",
      "ErrorMessage": "Rate exceeded for shard shardId-000000000001 in
        stream exampleStreamName under account 123456789."
    },
    {
      "SequenceNumber": "21269319989999637946712965403778482985",
      "ShardId": "shardId-000000000002"
    }
  ]
}
```

Which techniques will help mitigate this exception? (Select TWO.)

- A. Implement retries with exponential backoff
- B. Use a PutRecord API instead of PutRecords
- C. Reduce the frequency and/or size of the requests
- D. Use Amazon SNS instead of Kinesis.
- E. Reduce the number of KCL consumers.

Answer: AC

NEW QUESTION 390

- (Exam Topic 3)

An application uses Amazon DynamoDB as its backend database The application experiences sudden spikes in traffic over the weekend and variable but predictable spikes during weekdays The capacity needs to be set to avoid throttling errors at all times How can this be accomplished cost-effectively?

- A. Use provisioned capacity with AWS Auto Scaling throughout the week.
- B. Use on-demand capacity for the weekend and provisioned capacity with AWS Auto Scaling during the weekdays
- C. Use on-demand capacity throughout the week
- D. Use provisioned capacity with AWS Auto Scaling enabled during the weekend and reserved capacity enabled during the weekdays

Answer: A

NEW QUESTION 391

- (Exam Topic 3)

An application contains two components one component to handle HI IP 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 392

- (Exam Topic 3)

A developer implemented a static website hosted in Amazon S3 that makes web service requests hosted in Amazon API Gateway 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-Request-Headers header to the request

Answer: B

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-cors-console.html>

NEW QUESTION 394

- (Exam Topic 3)

A developer is creating an event handling system. To handle messages asynchronously, the developer created a standard Amazon SQS queue. Quality assurance testing reveals that some events were processed multiple times. What is the recommended way to ensure the events are not processed more than once?

- A. Change long polling to short polling.
- B. Use a FIFO queue and configure deduplication
- C. Convert the standard SQS queue into a FIFO queue
- D. Send the messages with message timers

Answer: C

NEW QUESTION 398

- (Exam Topic 3)

A company has deployed an application on AWS Elastic Beanstalk. The company has configured the Auto Scaling group that is associated with the Elastic Beanstalk environment to have five Amazon EC2 instances. If the capacity is fewer than four EC2 instances during the deployment, application performance degrades. The company is using the all-at-once deployment policy. What is the MOST cost-effective way to solve the deployment issue?

- A. Change the Auto Scaling group to six desired instances
- B. Change the deployment policy to traffic splitting. Specify an evaluation time of 1 hour.
- C. Change the deployment policy to rolling with additional batch. Specify a batch size of 1
- D. Change the deployment policy to rolling. Specify a batch size of 2.

Answer: C

NEW QUESTION 401

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 402

- (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 403

- (Exam Topic 3)

A developer must build a mobile application that allows users to read and write data from an Amazon DynamoDB table to store user state for each unique user. The solution needs to limit data access to allow users access only to their own data. Which solution below is the most secure?

- A. Embed AWS access credentials into the application and create DynamoDB queries that limit user access.
- B. Use Amazon Cognito identity pools to assign unique identifiers and provide user access.
- C. Modify the DynamoDB table to allow public read and writes, then add client-side filtering.
- D. Create a web portal for users to create an account on AWS Directory Service.

Answer: C

NEW QUESTION 404

- (Exam Topic 3)

A company is using an AWS Lambda function to process records from an Amazon Kinesis data stream. The company recently observed slow processing of the records. A developer notices that the iterator age metric for the function is increasing and that the Lambda run duration is constantly above normal. Which actions should the developer take to increase the processing speed? (Select TWO.)

- A. Increase the number of shards of the Kinesis data stream.
- B. Decrease the timeout of the Lambda function.
- C. Increase the memory that is allocated to the Lambda function.
- D. Decrease the number of shards of the Kinesis data stream.
- E. Increase the timeout of the Lambda function.

Answer: DE

NEW QUESTION 408

- (Exam Topic 3)

A company wants to make sure that only one user from its Admin group has the permanent right to delete an Amazon EC2 resource. There should be no changes in the existing policy under the Admin group. What should a developer use to meet these requirements?

- A. AWS managed policy
- B. Inline policy
- C. IAM trust relationship
- D. AWS Security Token Service (AWS STS)

Answer: B

NEW QUESTION 412

- (Exam Topic 3)

An application runs on multiple EC2 instances behind an ELB. Where is the session data best written so that it can be served reliably across multiple requests?

- A. Write data to Amazon ElastiCache.
- B. Write data to Amazon Elastic Block Store.
- C. Write data to Amazon EC2 instance Block Store.
- D. Write data to the root filesystem.

Answer: A

NEW QUESTION 414

- (Exam Topic 3)

A company is building a compute-intensive application that will run on a fleet of Amazon EC2 instances. The application uses attached Amazon EBS disks for storing data. The application will process sensitive information and all the data must be encrypted. What should a developer do to ensure the data is encrypted on disk without impacting performance?

- A. Configure the Amazon EC2 instance fleet to use encrypted EBS volumes for storing data.
- B. Add logic to write all data to an encrypted Amazon S3 bucket.
- C. Add a custom encryption algorithm to the application that will encrypt and decrypt all data.
- D. Create a new Amazon Machine Image (AMI) with an encrypted root volume and store the data on ephemeral disks.

Answer: A

NEW QUESTION 419

- (Exam Topic 4)

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

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

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

Answer: B

NEW QUESTION 423

- (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 426

- (Exam Topic 4)

A developer needs to deploy an application to AWS Elastic Beanstalk for a company. The application consists of a single Docker image. The company's automated continuous integration and continuous delivery (CI/CD) process builds the Docker image and pushes the image to a public Docker registry.

How should the developer deploy the application to Elastic Beanstalk?

- A. Create a Dockerfil
- B. Configure Elastic Beanstalk to build the application as a Docker image.
- C. Create a docker-compose.yml fil
- D. Use the Elastic Beanstalk CLI to deploy the application.
- E. Create a .zip file that contains the Docker imag
- F. Upload the .zip file to Elastic Beanstalk.
- G. Create a Dockerfil
- H. Run the Elastic Beanstalk CLI eb local run command in the same directory.

Answer: B

Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/docker.html#single-container-docker.deploy-remote> Deploy a remote Docker image to Elastic Beanstalk After testing your container locally, deploy it to an Elastic Beanstalk environment. Elastic Beanstalk uses the docker-compose.yml file to pull and run your image if you are using Docker Compose. Otherwise, Elastic Beanstalk uses the Dockerrun.aws.json instead.

Use the EB CLI to create an environment and deploy your image.

```
~/remote-docker$ eb create environment-name
```

NEW QUESTION 430

- (Exam Topic 4)

A developer has created a Java application that makes HTTP requests directly to AWS services. Application logging shows 5xx HTTP response codes that occur irregular intervals. The errors are affecting users.

How should the developer update the application to improve the application's resiliency?

- A. Revise the request content in the application code.
- B. Use the AWS SDK for Java to interact with AWS APIs.
- C. Scale out the application so that more instances of the application are running.
- D. Add additional logging to the application code.

Answer: B

NEW QUESTION 433

- (Exam Topic 4)

What are the MINIMUM properties required in the resources section of the AppSpace file for CodeDeploy to deploy the ECS service successfully?

- A. name, alias, currentversion, and targetversion
- B. TaskDefinition, ContainerName, and PlatformVersion
- C. TaskDefinitionContainerName, ContainerPort
- D. name, Currentversion, NetworkConfiguration, and Platform Version

Answer: A

NEW QUESTION 437

- (Exam Topic 4)

A company is migrating a web application from on premises to AWS. The company needs to move session storage from the application code to a shared service as part of the migration. The session storage data must be encrypted at rest.

Which AWS services meet these requirements? (Choose two.)

- A. Amazon ElastiCache for Redis
- B. Amazon ElastiCache for Memcached
- C. Amazon CloudWatch
- D. AWS CloudTrail
- E. Amazon DynamoDB

Answer: AE

Explanation:

<https://aws.amazon.com/blogs/security/amazon-elasticache-now-supports-encryption-for-elasticache-for-redis/>

A) ElastiCache for Redis is always a good option as a distributed cache for session management - <https://aws.amazon.com/getting-started/hands-on/building-fast-session-caching-with-amazon-elasticache-for-red>

It also supports encrypt at rest - <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/at-rest-encryption.html>

E) DynamoDB is also common to store session state with TTL support. And all user data stored in Amazon DynamoDB is fully encrypted at rest - <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/EncryptionAtRest.html>

NEW QUESTION 439

.....

Relate Links

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

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

Contact us

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

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