

## Amazon.DVA-C01.v2023-01-05.q214

Exam Code:	DVA-C01
Exam Name:	AWS Certified Developer Associate Exam
Certification Provider:	Amazon
Free Question Number:	214
Version:	v2023-01-05
# of views:	2227
# of Questions views:	2140
<a href="https://www.dumpsdb.com/dumps/Amazon/DVA-C01/Amazon.DVA-C01.v2023-01-05.q214">https://www.dumpsdb.com/dumps/Amazon/DVA-C01/Amazon.DVA-C01.v2023-01-05.q214</a>	

### NEW QUESTION: 1

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

What deployment method should be used?

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

**Answer: D (LEAVE A REPLY)**

Reference:

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

### NEW QUESTION: 2

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

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

**Answer: A (LEAVE A REPLY)**

### NEW QUESTION: 3

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

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

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

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 4

A company hosts a monolithic application on Amazon EC2 instances. The company starts converting some features of the application to a serverless architecture by using Amazon API Gateway and AWS Lambda. After the migration, some users report problems with payment processing. Upon inspection, a developer discovers that the Lambda function that calls the external payment API is taking longer than expected. Therefore, the API Gateway requests are timing out. What should the developer do to resolve this issue in the serverless architecture?

- A. Increase the API Gateway timeout duration to match the payment API time.
- B. Use Amazon Simple Queue Service (Amazon SQS) with API Gateway and the Lambda function to asynchronously call the payment API.
- C. Increase the Lambda function's memory to increase the network bandwidth and increase the speed of the payment API calls.
- D. Use the EC2 instances to make the API calls to the payment API.

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 5

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 (LEAVE A REPLY)**

<https://acloud.guru/forums/aws-certified-developer-associate/discussion/-KGngHzVQ03OpeAA9jSP/i-cant-answer-a-sample-question-pretty-worried-about-the-real-thing>

[https://acloud.guru/forums/aws-certified-developer-associate/discussion/-K5WKXRAIJdOu58GREF\\_/s3-question](https://acloud.guru/forums/aws-certified-developer-associate/discussion/-K5WKXRAIJdOu58GREF_/s3-question)

#### NEW QUESTION: 6

A Developer wants to debug an application by searching and filtering log data. The application logs are stored in Amazon CloudWatch Logs. The Developer creates a new metric filter to count exceptions in the application logs. However, no results are returned from the logs.

What is the reason that no filtered results are being returned?

- A. A setup of the Amazon CloudWatch interface VPC endpoint is required for filtering the CloudWatch Logs in the VPC.
- B. CloudWatch Logs only publishes metric data for events that happen after the filter is created.
- C. The log group for CloudWatch Logs should be first streamed to Amazon Elasticsearch Service before metric filtering returns the results.

D. Metric data points for logs groups can be filtered only after they are exported to an Amazon S3 bucket

**Answer: B (LEAVE A REPLY)**

Reference:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/MonitoringLogData.html>

#### NEW QUESTION: 7

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: B,D (LEAVE A REPLY)**

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

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

#### NEW QUESTION: 8

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

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

A. Push callback data to Amazon CloudWatch as a custom metric and use the CloudWatch alerting mechanisms to alert System Administrators.

B. Push callback data to Amazon Kinesis Data Streams and invoke an AWS Lambda function that stores data in Amazon DynamoDB and sends the required alerts.

C. Push callback data to AWS X-Ray and use AWS Lambda to query, display, and alert on exceeded thresholds.

D. Push callback data to an Amazon RDS database that can be queried to show historical data and to alert on exceeded thresholds.

**Answer: (SHOW ANSWER)**

#### NEW QUESTION: 9

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

A. HTTP

B. UDP

C. SMS

D. DynamoDB

E. Named Pipes

**Answer: A,C (LEAVE A REPLY)**

Reference:

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

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

In order for customers to have broad flexibility of delivery mechanisms, Amazon SNS supports notifications over multiple transport protocols.

Customers can select one the following transports as part of the subscription requests:

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

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

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

"SMS" - Messages are sent to registered phone numbers as SMS text messages.

#### **NEW QUESTION: 10**

A company is developing a new web application in Python A developer must deploy the application using AWS Elastic Beanstalk from the AWS Management Console The developer creates an Elastic Beanstalk source bundle to upload using the console Which of the following are requirements when creating the source bundle? (Select TWO.)

- A. The source bundle must be compressed with any required dependencies in a top-level parent folder
- B. The source bundle must not include a top-level directory.
- C. The source bundle must be created as a single zip or war file
- D. The source bundle must include the ebextensions.yaml file.
- E. The source bundle must be uploaded into Amazon EFS.

**Answer: B,C ([LEAVE A REPLY](#))**

#### **NEW QUESTION: 11**

An ecommerce application is using Amazon Simple Notification Service (Amazon SNS) with an AWS Lambda subscription to save all new orders into an Amazon DynamoDB table The company wants to record all the orders that are more than a certain amount of money in a separate table The company wants to avoid changes to the processes that post orders to Amazon SNS or the current Lambda function that saves the orders to the DynamoDB table How can a developer implement this feature with the LEAST change to the existing application?

- A. Modify the Lambda code to filter the orders and save the appropriate orders to a separate table
- B. Create another SNS topic, and also send orders in that topic Create a Lambda subscription with a numeric value filter option to save the appropriate orders to a separate table
- C. Create another Lambda subscription with the SNS message attribute value matching a filter option to save the appropriate orders to a separate table
- D. Create another Lambda subscription with the SNS message numeric value matching a filter option to save the appropriate orders to a separate table

**Answer: ([SHOW ANSWER](#))**

#### **NEW QUESTION: 12**

What type of block cipher does Amazon S3 offer for server side encryption?

- A. Blowfish
- B. Advanced Encryption Standard
- C. Triple DES
- D. RC5

**Answer: B (LEAVE A REPLY)**

### NEW QUESTION: 13

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

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

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

**Answer: (SHOW ANSWER)**

Explanation

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

### NEW QUESTION: 14

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

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

**Answer: A (LEAVE A REPLY)**

Explanation

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

### NEW QUESTION: 15

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 (LEAVE A REPLY)**

<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: 16

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

Which of the following steps could resolve the issue?

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

**Answer:** ([SHOW ANSWER](#))

[https://docs.aws.amazon.com/vpc/latest/userguide/VPC\\_NAT\\_Instance.html#NATInstance](https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Instance.html#NATInstance)

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (610 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### **NEW QUESTION: 17**

A Developer is storing sensitive documents in Amazon S3 that will require encryption at rest. The encryption keys must be rotated annually, at least. What is the easiest way to achieve this?

- A. Use AWS KMS with automatic key rotation
- B. Encrypt the data before sending it to Amazon S3
- C. Import a custom key into AWS KMS with annual rotation enabled
- D. Export a key from AWS KMS to encrypt the data

**Answer:** C ([LEAVE A REPLY](#))

#### **NEW QUESTION: 18**

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 ([LEAVE A REPLY](#))

Reference:

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

#### **NEW QUESTION: 19**

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

Which solution will meet this requirement?

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

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 20**

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

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 21**

An application is running on an EC2 instance. The Developer wants to store an application metric in Amazon CloudWatch.

What is the best practice for implementing this requirement?

- A.** Use the PUT Object API call to send data to an S3 bucket. Use an event notification to invoke a Lambda function to publish data to CloudWatch.
- B.** Publish the metric data to an Amazon Kinesis Stream using a PutRecord API call. Subscribe a Lambda function that publishes data to CloudWatch.
- C.** Use the CloudWatch PutMetricData API call to submit a custom metric to CloudWatch. Provide the required credentials to enable the API call.
- D.** Use the CloudWatch PutMetricData API call to submit a custom metric to CloudWatch. Launch the EC2 instance with the required IAM role to enable the API call.

**Answer:** ([SHOW ANSWER](#))

Reference:

[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_use\\_switch-role-ec2.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html)

#### **NEW QUESTION: 22**

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:** ([SHOW ANSWER](#))

Reference:

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

#### **NEW QUESTION: 23**

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. Modify the ECS task definition to increase the task memory
- C. Scale the ECS cluster to contain more ECS instances
- D. Add read capacity units (RCUs) to the DB instance

**Answer:** A ([LEAVE A REPLY](#))

#### **NEW QUESTION: 24**

A company is expanding the compatibility of its photo-sharing mobile app to hundreds of additional devices with unique screen dimensions and resolutions. Photos are stored in Amazon S3 in their original format and resolution. The company uses an Amazon CloudFront distribution to serve the photos. The app includes the dimension and resolution of the display as GET parameters with every request.

A developer needs to implement a solution that optimizes the photos that are served to each device to reduce load time and increase photo quality. Which solution will meet these requirements MOST cost-effectively?

- A. Create a Lambda@Edge function that optimizes the photos upon request and returns the photos as a response. In the same function, store a copy of processed photos on Amazon S3 for subsequent requests.
- B. Use S3 Batch Operations to invoke an AWS Lambda function to create new variants of the photos with the required dimensions and resolutions. Create a dynamic CloudFront origin that automatically maps the request of each device to the corresponding photo variant.
- C. Use S3 Batch Operations to invoke an AWS Lambda function to create new variants of the photos with the required dimensions and resolutions. Create a Lambda@Edge function to route requests to the corresponding photo variant by using request headers.
- D. Create a Lambda@Edge function that optimizes the photos upon request and returns the photos as a response. Change the CloudFront TTL cache policy to the maximum value possible.

**Answer:** B ([LEAVE A REPLY](#))

#### **NEW QUESTION: 25**

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

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

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

**Answer: B (LEAVE A REPLY)**

[https://docs.aws.amazon.com/en\\_us/amazondynamodb/latest/developerguide/Streams.Lambda.Tutorial.html](https://docs.aws.amazon.com/en_us/amazondynamodb/latest/developerguide/Streams.Lambda.Tutorial.html)

Create an event source mapping to tell Lambda to send records from your stream to a Lambda function. You can create multiple event source mappings to process the same data with multiple Lambda functions, or process items from multiple streams with a single function.

#### **NEW QUESTION: 26**

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: (SHOW ANSWER)**

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

#### **NEW QUESTION: 27**

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

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

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

**Answer: (SHOW ANSWER)**

#### **NEW QUESTION: 28**

A Developer is creating an Auto Scaling group whose instances need to publish a custom metric to Amazon CloudWatch.

Which method would be the MOST secure way to authenticate a CloudWatch PUT request?

- A.** Create an IAM user with PutMetricData permission, and modify the Auto Scaling launch configuration to inject the user credentials into the instance user data.
- B.** Create an IAM user with PutMetricData permission and put the user credentials in a private repository; have applications pull the credentials as needed.
- C.** Create an IAM role with PutMetricData permission and modify the Auto Scaling launching configuration to launch instances using that role.
- D.** Modify the CloudWatch metric policies to allow the PutMetricData permission to instances from the Auto Scaling group.

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 29**

A company is providing read access to objects in an Amazon S3 bucket for different customers. The company uses IAM permissions to restrict access to the S3 bucket. The customers can access only their own files. Due to a regulation requirement, the company needs to enforce encryption in transit for interactions with Amazon S3. Which solution will meet these requirements?

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

Answer: [\(SHOW ANSWER\)](#)

**NEW QUESTION: 30**

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

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

Answer: **B** [\(LEAVE A REPLY\)](#)

**NEW QUESTION: 31**

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

What could be causing this issue?

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

Answer: **A** [\(LEAVE A REPLY\)](#)

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (610 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

**NEW QUESTION: 32**

A developer has launched an application that calls an API by way of Amazon API Gateway. It offers information that changes several times a day, but is not updated in real time. The application has become so popular that the API endpoint is overloaded and that traffic to the endpoint must be reduced.

What can the developer do to address the performance issues?

- A. Enable an Auto Scaling group on the endpoint service and database.
- B. Enable API caching in Amazon ElastiCache.
- C. Create an additional API Gateway and use an Application Load Balancer

**Answer:** ([SHOW ANSWER](#))

#### NEW QUESTION: 33

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

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

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

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

**Answer:** A ([LEAVE A REPLY](#))

#### NEW QUESTION: 34

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

What is the MOST likely cause of this error?

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

**Answer:** ([SHOW ANSWER](#))

#### NEW QUESTION: 35

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

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

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 36**

A web application is using Amazon Kinesis Streams for clickstream data that may not be consumed for up to 12 hours. How can the Developer implement encryption at rest for data within the Kinesis Streams?

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

**Answer:** ([SHOW ANSWER](#))

Reference:

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

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

#### **NEW QUESTION: 37**

An existing serverless application processes uploaded image files. The process currently uses a single Lambda function that takes an image file, performs the processing, and stores the file in Amazon S3. Users of the application now require thumbnail generation of the images. Users want to avoid any impact to the time it takes to perform the image uploads.

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

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

**Answer:** **C** ([LEAVE A REPLY](#))

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

#### **NEW QUESTION: 38**

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

What steps must be followed?

- A. Invoke the Encrypt API passing the plaintext data that must be encrypted, then reference the customer managed key ARN in the KeyId parameter

- B. Invoke the GenerateRandom API to get a data encryption key, then use the data encryption key to encrypt the data
- C. Invoke the GenerateDataKey API to retrieve the encrypted version of the data encryption key to encrypt the data
- D. Invoke the GenerateDataKey API to retrieve the plaintext version of the data encryption key to encrypt the data

**Answer: (SHOW ANSWER)**

Reference:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingClientSideEncryption.html> GenerateDataKey API: Generates a unique data key. This operation returns a plaintext copy of the data key and a copy that is encrypted under a customer master key (CMK) that you specify. You can use the plaintext key to encrypt your data outside of KMS and store the encrypted data key with the encrypted data.

#### **NEW QUESTION: 39**

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

**Answer: A,D (LEAVE A REPLY)**

#### **NEW QUESTION: 40**

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 (LEAVE A REPLY)**

Explanation

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

#### **NEW QUESTION: 41**

A Developer must repeatedly and consistently deploy a serverless RESTful API on AWS.

Which techniques will work? (Choose two.)

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

**Answer: (SHOW ANSWER)**

Explanation

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

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

#### **NEW QUESTION: 42**

When a Simple Queue Service message triggers a task that takes 5 minutes to complete, which process below will result in successful processing of the message and remove it from the queue while minimizing the chances of duplicate processing?

- A. Retrieve the message with an increased visibility timeout, process the message, delete the message from the queue
- B. Retrieve the message with an increased visibility timeout, delete the message from the queue, process the message
- C. Retrieve the message with increased DelaySeconds, process the message, delete the message from the queue
- D. Retrieve the message with increased DelaySeconds, delete the message from the queue, process the message

**Answer: A (LEAVE A REPLY)**

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

#### **NEW QUESTION: 43**

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

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

**Answer: C (LEAVE A REPLY)**

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

#### **NEW QUESTION: 44**

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

**Answer: A (LEAVE A REPLY)**

#### **NEW QUESTION: 45**

A Developer is creating a mobile application that will not require users to log in.

What is the MOST efficient method to grant users access to AWS resources?

- A. Use an identity provider to securely authenticate with the application.
- B. Create an AWS Lambda function to create an IAM user when a user accesses the application.

- C. Create credentials using AWS KMS and apply these credentials to users when using the application.
- D. Use Amazon Cognito to associate unauthenticated users with an IAM role that has limited access to resources.

**Answer: D (LEAVE A REPLY)**

Reference:

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

#### **NEW QUESTION: 46**

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

How can the developer authenticate and authorize users?

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

**Answer: (SHOW ANSWER)**

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (610 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### **NEW QUESTION: 47**

A Developer decides to store highly secure data in Amazon S3 and wants to implement server-side encryption (SSE) with granular control of who can access the master key. Company policy requires that the master key be created, rotated, and disabled easily when needed, all for security reasons.

Which solution should be used to meet these requirements?

- A. SSE with Amazon S3 managed keys (SSE-S3)
- B. SSE with customer provided encryption keys
- C. SSE with AWS Secrets Manager
- D. SSE with AWS KMS managed keys (SSE KMS)

**Answer: D (LEAVE A REPLY)**

#### **NEW QUESTION: 48**

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

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

**Answer: A** ([LEAVE A REPLY](#))

**NEW QUESTION: 49**

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

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

**Answer: C** ([LEAVE A REPLY](#))

**NEW QUESTION: 50**

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

**Answer: A** ([LEAVE A REPLY](#))

**NEW QUESTION: 51**

A company is providing read access to objects in an Amazon S3 bucket for different customers The company uses IAM permissions to restrict access to the S3 bucket The customers can access only their own files Due to a regulation requirement the company needs to enforce encryption in transit for interactions with Amazon S3 Which solution will meet these requirements'?

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

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 52**

Given the following AWS CloudFormation template:



What is the MOST efficient way to reference the new Amazon S3 bucket from another AWS CloudFormation template?

- A. Create a custom AWS CloudFormation resource that gets the bucket name from the ContentBucket resource of the first stack.

- B. Use Fn::Include to include the existing template in other template and use the ContentBucket resource directly
- C. Add Exported: True to the ContentBucket in the original template and use ImportResource in other templates.
- D. Add an Export declaration to the outputs section of the original template and use ImportValue in other templates.

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 53**

A company is developing a serverless ecommerce web application. The application needs to make coordinated, all-or-nothing changes to multiple items in the company's inventory table in Amazon DynamoDB.

Which solution will meet these requirements?

- A. Enable transactions for the DynamoDB table. Use the BatchWriteItem operation to update the items.
- B. Use the TransactWriteItem operation to group the changes. Update the items in the table.
- C. Set up a FIFO queue using Amazon SQS. Group the changes in the queue. Update the table based on the grouped changes.
- D. Create a transaction table in an Amazon Aurora DB cluster to manage the transactions. Write a backend process to sync the Aurora DB table and the DynamoDB table.

**Answer:** ([SHOW ANSWER](#))

Reference:

TransactWriteItems is a synchronous write operation that groups up to 25 action requests. The BatchWriteItem operation puts or deletes multiple items in one or more tables. <https://docs.aws.amazon.com/AWSJavaScriptSDK/latest/AWS/DynamoDB.html>

#### **NEW QUESTION: 54**

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 a central FTP server to deploy the application to Elastic Beanstalk.
- B. Create an AWS CodeCommit repository, allow the Developers to commit code to it, and then directly deploy the code to Elastic Beanstalk.
- C. 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.
- D. Allow the Developers to upload the code to an Amazon S3 bucket, and deploy it directly to Elastic Beanstalk.

**Answer:** B ([LEAVE A REPLY](#))

#### **NEW QUESTION: 55**

An on-premises legacy application is caching data files locally and writing shared images to local disks.

What is necessary to allow for horizontal scaling when migrating the application to AWS?

- A. Modify the application to read and write cache data on Amazon S3, and also store shared images on S3.
- B. Modify the application to read and write cache data on Amazon S3, while continuing to write shared images to local disks.
- C. Modify the application to have both shared images and caching data written to Amazon EBS.
- D. Modify the application to use Amazon S3 for serving shared images; cache data can then be written to local disks.

**Answer:** D ([LEAVE A REPLY](#))

#### **NEW QUESTION: 56**

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

How should the developer accomplish this?

- A. After the new file is created in Lambda, use cURL to invoke the CodeCommit API. Send the file to the repository.
- B. When the Lambda function starts, use the Git CLI to Clone the repository. Check the new file into the cloned repository and push the change.
- C. Upload the new file to an Amazon S3 bucket. Create an AWS step Function to accept S3 events. In the step Function, add the new file to the repository.
- D. Use an AWS SDK to instantiate a CodeCommit Client. Invoke the put \_ file method to add the file to the repository.

**Answer: D (LEAVE A REPLY)**

#### **NEW QUESTION: 57**

An Amazon S3 bucket, "myawsbucket" is configured with website hosting in Tokyo region, what is the region-specific website endpoint?

- A. www.myawsbucket.ap-northeast-1.amazonaws.com
- B. myawsbucket.s3-website-ap-northeast-1.amazonaws.com
- C. myawsbucket.amazonaws.com
- D. myawsbucket.tokyo.amazonaws.com

**Answer: B (LEAVE A REPLY)**

Depending on your Region, your Amazon S3 website endpoint follows one of these two formats.

s3-website dash (-) Region - http://bucket-name.s3-website-Region.amazonaws.com

s3-website dot (.) Region - http://bucket-name.s3-website.Region.amazonaws.com

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

#### **NEW QUESTION: 58**

A Developer is building a three-tier web application that should be able to handle a minimum of 5000 requests per minute. Requirements state that the web tier should be completely stateless while the application maintains session state for the users.

How can session data be externalized, keeping latency at the LOWEST possible value?

- A. Create an Amazon RDS instance, then implement session handling at the application level to leverage a database inside the RDS database instance for session data storage
- B. Implement a shared file system solution across the underlying Amazon EC2 instances, then implement session handling at the application level to leverage the shared file system for session data storage
- C. Create an Amazon ElastiCache Memcached cluster, then implement session handling at the application level to leverage the cluster for session data storage
- D. Create an Amazon DynamoDB table, then implement session handling at the application level to leverage the table for session data storage

**Answer: (SHOW ANSWER)**

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

#### **NEW QUESTION: 59**

A company maintains a REST service using Amazon API Gateway and the API Gateway native API key validation. The company recently launched a new registration page, which allows users to sign up for the service. The registration page creates a new API key using CreateApiKey and sends the new key to the user.

When the user attempts to call the API using this key, the user receives a 403 Forbidden error. Existing users are unaffected and can still call the API.

What code updates will grant these new users access to the API?

- A. The createDeployment method must be called so the API can be redeployed to include the newly created API key.
- B. The updateAuthorizer method must be called to update the API's authorizer to include the newly created API key.
- C. The importApiKeys method must be called to import all newly created API keys into the current stage of the API.
- D. The createUsagePlanKey method must be called to associate the newly created API key with the correct usage plan.

**Answer: D (LEAVE A REPLY)**

Explanation

<https://stackoverflow.com/questions/39061041/using-an-api-key-in-amazon-api-gateway>

#### NEW QUESTION: 60

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

How can the Developer address this issue?

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

**Answer: (SHOW ANSWER)**

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

#### NEW QUESTION: 61

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: (SHOW ANSWER)**

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.

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (610 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### **NEW QUESTION: 62**

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

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

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

**Answer: A (LEAVE A REPLY)**

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

#### **NEW QUESTION: 63**

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

What is the fastest way to process all the files?

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

**Answer: C (LEAVE A REPLY)**

#### **NEW QUESTION: 64**

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

- A.** Create an AWS Lambda rotation function that has appropriate IAM permissions Store the password in AWS Secrets Manager Configure Secrets Manager to rotate the password by using the Lambda function
- B.** Encrypt the existing password with AWS Key Management Service (AWS KMS) Export the existing password Generate a random password with AWS KMS Use the AWS KMS password renewal feature to replace the existing password with the new password.

**C.** Integrate AWS Systems Manager Parameter Store with a Key Management Interoperability Protocol (KMIP)-compliant third-party secret manager to enable third-party database password rotation on AWS

**D.** Create an AWS Lambda rotation function that has appropriate IAM permissions Store the password in AWS Systems Manager Parameter Store Configure Parameter Store to rotate the password by using the Lambda function

**Answer: C (LEAVE A REPLY)**

#### **NEW QUESTION: 65**

An Amazon S3 bucket, "myawsbucket" is configured with website hosting in Tokyo region, what is the region-specific website endpoint?

**A.** myawsbucket.s3-website-ap-northeast-1.amazonaws.com

**B.** www.myawsbucket.ap-northeast-1.amazonaws.com

**C.** myawsbucket.amazonaws.com

**D.** myawsbucket.tokyo.amazonaws.com

**Answer: A (LEAVE A REPLY)**

#### **NEW QUESTION: 66**

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

How should the environment variables be passed to the container?

**A.** Define an array that includes the environment variables under the entrypoint parameter within the task definition

**B.** Define an array that includes the environment variables under the environment parameter within the service definition

**C.** Define in array that includes the environment variables under the entryPoint parameter within the service definition

**D.** Define an array that includes the environment variables under the environment parameter within the task definition

**Answer: (SHOW ANSWER)**

#### **NEW QUESTION: 67**

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

**A.** Integrate Amazon CloudFront with Amazon S3.

**B.** Enable Amazon S3 cross-region replication.

**C.** Delete expired Amazon S3 server log files.

**D.** Configure Amazon S3 lifecycle rules.

**E.** Randomize Amazon S3 key name prefixes.

**Answer: A,E (LEAVE A REPLY)**

Reference:

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: 68**

An Amazon RDS database instance is used by many applications to look up historical data. The query rate is relatively constant. When the historical data is updated each day, the resulting write traffic slows the read query performance and affects all application users.

What can be done to eliminate the performance impact on application users?

- A. Make sure Amazon RDS is Multi-AZ so it can better absorb increased traffic.
- B. Create an RDS Read Replica and direct all read traffic to the replica.
- C. Implement Amazon ElastiCache in front of Amazon RDS to buffer the write traffic.
- D. Use Amazon DynamoDB instead of Amazon RDS to buffer the read traffic.

**Answer: B (LEAVE A REPLY)**

Reference:

<https://aws.amazon.com/rds/details/read-replicas/>

#### **NEW QUESTION: 69**

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

**Answer: A (LEAVE A REPLY)**

#### **NEW QUESTION: 70**

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

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

**Answer: A (LEAVE A REPLY)**

#### **NEW QUESTION: 71**

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. Use server-side encryption with AWS KMS managed keys (SSE-KMS) and download the documents using HTTPS
- B. Create a presigned S3 URL using the AWS SDK with an expiration time of 15 minutes

- C. Copy the documents to a separate S3 bucket that has a lifecycle policy for deletion after 15 minutes
- D. Modify the S3 bucket policy to only allow specific users to download the documents Revert the change after 15 minutes.

**Answer: B (LEAVE A REPLY)**

#### **NEW QUESTION: 72**

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 (LEAVE A REPLY)**

Reference:

<https://stackoverflow.com/questions/39582752/do-global-secondary-index-gsi-in-dynamodb-impact-tables-provision-capacity>

<https://medium.com/@synchronphoto/amazon-dynamodb-provisioning-write-capacity-for-global-secondary-indexes-gsis-eaa4d4dce1df>

#### **NEW QUESTION: 73**

A Development team has pushed out 10 applications running on several Amazon EC2 instances. The Operations team is asking for a graphical representation of one key performance metric for each application.

These metrics should be available on one screen for easy monitoring.

Which steps should the Developer take to accomplish this using Amazon CloudWatch?

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

**Answer: A (LEAVE A REPLY)**

#### **NEW QUESTION: 74**

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

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

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

**Answer: D (LEAVE A REPLY)**

#### **NEW QUESTION: 75**



**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (610 Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

#### **NEW QUESTION: 77**

A Development team is working on a case management solution that allows medical claims to be processed and reviewed. Users log in to provide information related to their medical and financial situations.

As part of the application, sensitive documents such as medical records, medical imaging, bank statements, and receipts are uploaded to Amazon S3. All documents must be securely transmitted and stored. All access to the documents must be recorded for auditing.

What is the MOST secure approach?

- A. Use S3 default encryption using Advanced Encryption Standard-256 (AES-256) on the destination bucket.
- B. Use Amazon Cognito for authorization and authentication to ensure the security of the application and documents.
- C. Use AWS Lambda to encrypt and decrypt objects as they are placed into the S3 bucket.
- D. Use client-side encryption/decryption with Amazon S3 and AWS KMS.

**Answer: D (LEAVE A REPLY)**

Explanation

[https://aws.amazon.com/s3/faqs/?nc1=h\\_ls](https://aws.amazon.com/s3/faqs/?nc1=h_ls).

#### **NEW QUESTION: 78**

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 duster.

How should the application authenticate with AWS services in production?

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

**Answer: A (LEAVE A REPLY)**

#### **NEW QUESTION: 79**

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 instance role does not have "xray:PutTraceSegments" and "xray:PutTelemetryRecords" permissions.
- C. The instance role does not have "xray:BatchGetTraces" and "xray:GetTraceGraph" permissions.
- D. The X-Ray endpoint specified in the application configuration is incorrect.
- E. The X-Ray daemon is not installed on the EC2 instance.

**Answer: B,E (LEAVE A REPLY)**

**NEW QUESTION: 80**

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

What is the reason for the duplicates?

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

**Answer: A (LEAVE A REPLY)**

Reference:

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

**NEW QUESTION: 81**

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

What solution would give the BEST performance increase?

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

**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 82**

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 (LEAVE A REPLY)**

Explanation

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

**NEW QUESTION: 83**

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. Update and add AWS X-ray daemon code to relevant parts of the Lambda function to set up the trace.
- C. Trigger a Lambda function from the application logs in Amazon CloudWatch to submit tracing data to AWS X-Ray
- D. Use an IAM execution role to give the Lambda function permissions and enabled tracing.

**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 84**

A company maintains a REST service using Amazon API Gateway and the API Gateway native API key validation. The company recently launched a new registration page, which allows users to sign up for the service. The registration page creates a new API key using CreateApiKey and sends the new key to the user.

When the user attempts to call the API using this key, the user receives a 403 Forbidden error. Existing users are unaffected and can still call the API. What code updates will grant these new users access to the API?

- A. The createUsagePlanKey method must be called to associate the newly created API key with the correct usage plan.
- B. The updateAuthorizer method must be called to update the API's authorizer to include the newly created API key.
- C. The importApiKeys method must be called to import all newly created API keys into the current stage of the API.
- D. The createDeployment method must be called so the API can be redeployed to include the newly created API key.

**Answer: C (LEAVE A REPLY)**

**NEW QUESTION: 85**

A company is building an application to track athlete performance using an Amazon DynamoDB table. Each item in the table is identified by a partition key (user\_id) and a sort key (sport\_name). The table design is shown below:

(Note: Not all table attributes are shown)

A Developer is asked to write a leaderboard application to display the top performers (user\_id) based on the score for each sport\_name.

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

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

**Answer: (SHOW ANSWER)**

Explanation

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

[https://docs.aws.amazon.com/zh\\_cn/amazondynamodb/latest/developerguide/GSI.html](https://docs.aws.amazon.com/zh_cn/amazondynamodb/latest/developerguide/GSI.html)

**NEW QUESTION: 86**

An AWS Lambda function must access an external site by using a regularly rotated user name and password.

These items must be kept securely and cannot be stored in the function code.

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

- A. AWS Certificate Manager (ACM)
- B. AWS Systems Manager Parameter Store
- C. AWS KMS
- D. AWS Trusted Advisor
- E. Amazon GuardDuty

**Answer: B,C (LEAVE A REPLY)**

**NEW QUESTION: 87**

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

**Answer:** [\(SHOW ANSWER\)](#)

#### **NEW QUESTION: 88**

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

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

**Answer:** [B \(LEAVE A REPLY\)](#)

#### **NEW QUESTION: 89**

What type of block cipher does Amazon S3 offer for server side encryption?

- A. Triple DES
- B. Advanced Encryption Standard
- C. Blowfish
- D. RC5

**Answer:** [B \(LEAVE A REPLY\)](#)

Explanation

<https://docs.aws.amazon.com/AmazonS3/latest/dev/serv-side-encryption.html>

#### **NEW QUESTION: 90**

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

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

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

**Answer:** [B \(LEAVE A REPLY\)](#)

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

#### **NEW QUESTION: 91**

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. DownloadBundle
- B. BeforeInstall
- C. AfterInstall
- D. ValtdateService

**Answer:** ([SHOW ANSWER](#))

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (**610 Q&As Dumps, 40%OFF** **Special Discount: Exam-Tests**)

#### **NEW QUESTION: 92**

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

How can a developer meet these requirements'?

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

**Answer: C** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 93**

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

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

**Answer: A,C,D,F** ([LEAVE A REPLY](#))

Explanation

Physical and Environmental Security

AWS's data centers are state of the art, utilizing innovative architectural and engineering approaches.

Amazon has many years of experience in designing, constructing, and operating large-scale data centers. This experience has been applied to the AWS platform and infrastructure. AWS data centers are housed in nondescript facilities. Physical access is strictly controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized

staff must pass two-factor authentication a minimum of two times to access data center floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff.



When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals.



AWS uses the techniques detailed in DoD 5220.22-M (National Industrial Security Program Operating Manual) or NIST 800-88 (Guidelines for Media Sanitization) to destroy data as part of the decommissioning process.



All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.  
Storage Decommissioning

#### **NEW QUESTION: 94**

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

**Answer: A** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 95**

A Developer created a Lambda function for a web application backend. When testing the Lambda function from the AWS Lambda console, the Developer can see that the function is being executed, but there is no log data being generated in Amazon CloudWatch Logs, even after several minutes.

What could cause this situation?

- A. The Lambda function is missing a target CloudWatch Log group.
- B. The Lambda function is missing CloudWatch Logs as a source trigger to send log data.
- C. The execution role for the Lambda function is missing permissions to write log data to the CloudWatch Logs.
- D. The Lambda function does not have any explicit log statements for the log data to send it to CloudWatch Logs.

**Answer: C** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 96**

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

- A. DescribeInstances
- B. DescribeAMIs
- C. DescribeImages

D. GetAMIs

E. You cannot retrieve a list of AMIs as there are over 10,000 AMIs

**Answer: C (LEAVE A REPLY)**

[https://docs.aws.amazon.com/AWSEC2/latest/APIReference/API\\_DescribeImages.html](https://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_DescribeImages.html)

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

#### **NEW QUESTION: 97**

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. Trigger the Lambda function synchronously from the SNS message.

B. Enable a DynamoDB stream that publishes an SNS message. Trigger the Lambda function asynchronously from the SNS message.

C. Enable a DynamoDB stream, and trigger the Lambda function synchronously from the stream.

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

**Answer: C (LEAVE A REPLY)**

Reference:

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

#### **NEW QUESTION: 98**

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. Call the CreatePlatformEndPoint API function to register multiple device tokens.

B. Bulk upload the device tokens contained in a CSV file via the AWS Management Console.

C. Implement a token vending service to handle the registration.

D. Let the push notification service (e.g. Amazon Device Messaging) handle the registration.

**Answer: (SHOW ANSWER)**

#### **NEW QUESTION: 99**

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 CodeCommit. Create a CodePipeline to automate unit testing. Use Amazon CloudWatch to trigger notification of failure events.

B. Store the source code in AWS CodeCommit. Create a CodePipeline to automate unit testing. Use Amazon SNS to trigger notifications of failure events.

C. Store the source code in GitHub. Create a CodePipeline to automate unit testing. Use Amazon SES to trigger notifications of failure events.

D. Store the source code on GitHub. Create a CodePipeline to automate unit testing. Use Amazon CloudWatch to trigger notifications of failure events.

**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 100**

A company is using AWS Elastic Beanstalk to deploy a three-tier application. The application uses an Amazon RDS DB instance as the database tier. The company wants to decouple the DB instance from the Elastic Beanstalk environment.

Which combination of steps should a developer take to meet this requirement? (Select TWO.)

- A. Use the AWS CLI to decouple the DB instance.
- B. Create a new DB instance from a snapshot of the previous DB instance.
- C. Modify the current Elastic Beanstalk environment to connect to the DB instance.
- D. Create a new Elastic Beanstalk environment that connects to the DB instance.
- E. Use the Elastic Beanstalk CLI to decouple the DB instance.

**Answer: B,C ([LEAVE A REPLY](#))**

**NEW QUESTION: 101**

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

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

- A. AWS Certificate Manager (ACM)
- B. AWS KMS
- C. Amazon GuardDuty
- D. AWS Systems Manager Parameter Store
- E. AWS Trusted Advisor

**Answer: ([SHOW ANSWER](#))**

**NEW QUESTION: 102**

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 ([LEAVE A REPLY](#))**

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

<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: 103**

A Developer is writing a serverless application that requires that an AWS Lambda function be invoked every 10 minutes.

What is an automated and serverless way to trigger the function?

- A. Deploy an Amazon EC2 instance based on Linux, and edit its `/etc/crontab` file by adding a command to periodically invoke the Lambda function.
- B. Configure an environment variable named `PERIOD` for the Lambda function. Set the value to 600.

C. Create an Amazon CloudWatch Events rule that triggers on a regular schedule to invoke the Lambda function.

D. Create an Amazon SNS topic that has a subscription to the Lambda function with a 600-second timer.

**Answer: C (LEAVE A REPLY)**

Reference:

<https://aws.amazon.com/blogs/architecture/a-serverless-solution-for-invoking-aws-lambda-at-a-sub-minute-frequency/>

#### NEW QUESTION: 104

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. Update the AWS codeBuild build specification to include a phase for running unit tests.

B. Install the AWS CodeDeploy agent on an Amazon EC2 instance to run unit tests.

C. Create a testing branch in AWS CodeCommit to run unit tests.

D. Create a separate codePipeline pipeline to run unit tests.

**Answer: A (LEAVE A REPLY)**

#### NEW QUESTION: 105

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

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

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

A. "Mapping" : "AWS::CloudFormation::NestedStack",

B. "Mapping" : "AWS::CloudFormation::Stack",

C. "Type" : "AWS::CloudFormation::NestedStack",

D. "Type" : "AWS::CloudFormation::Stack",

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 106

An application reads data from an Amazon DynamoDB table. Several times a day, for a period of 15 seconds, the application receives multiple ProvisionedThroughputExceeded errors.

How should this exception be handled?

- A. Create a new global secondary index for the table to help with the additional requests.
- B. Retry the failed read requests with exponential backoff.
- C. Immediately retry the failed read requests.
- D. Use the DynamoDB "UpdateItem" API to increase the provisioned throughput capacity of the table.

**Answer: B (LEAVE A REPLY)**

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

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (610 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### **NEW QUESTION: 107**

A development team is designing a mobile app that requires multi-factor authentication

Which steps should be taken to achieve this? (Select TWO)

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

**Answer: B,E (LEAVE A REPLY)**

#### **NEW QUESTION: 108**

A developer has written an Amazon kinesis Data streams application. As usage grows and traffic over time, the application is regularly receiving provisionedThroughputExceededException error messages.

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

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

**Answer: (SHOW ANSWER)**

Reference:

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

#### **NEW QUESTION: 109**

A Developer created a Lambda function for a web application backend. When testing the Lambda function from the AWS Lambda console, the Developer can see that the function is being executed, but there is no log data being generated in Amazon CloudWatch Logs, even after several minutes.

What could cause this situation?

- A. The Lambda function does not have any explicit log statements for the log data to send it to CloudWatch Logs.
- B. The Lambda function is missing CloudWatch Logs as a source trigger to send log data.
- C. The execution role for the Lambda function is missing permissions to write log data to the CloudWatch Logs.
- D. The Lambda function is missing a target CloudWatch Log group.

**Answer: C (LEAVE A REPLY)**

Reference:

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

#### **NEW QUESTION: 110**

A Developer has a stateful web server on-premises that is being migrated to AWS. The Developer must have greater elasticity in the new design. How should the Developer re-factor the application to make it more elastic? (Choose two.)

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

**Answer: B,E (LEAVE A REPLY)**

#### **NEW QUESTION: 111**

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

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

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

**Answer: C (LEAVE A REPLY)**

#### **NEW QUESTION: 112**

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

**Answer: (SHOW ANSWER)**

#### **NEW QUESTION: 113**

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

**Answer: A (LEAVE A REPLY)**

#### NEW QUESTION: 114

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

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

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

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 115

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

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 116

A legacy service has an XML-based SOAP interface. The Developer wants to expose the functionality of the service to external clients with the Amazon API Gateway. Which technique will accomplish this?

- A. Create a RESTful API with the API Gateway; transform the incoming JSON into a valid XML message for the SOAP interface using mapping templates.
- B. Create a RESTful API with the API Gateway; pass the incoming JSON to the SOAP interface through an Application Load Balancer.
- C. Create a RESTful API with the API Gateway; pass the incoming XML to the SOAP interface through an Application Load Balancer.
- D. Create a RESTful API with the API Gateway; transform the incoming XML into a valid message for the SOAP interface using mapping templates.

**Answer: A (LEAVE A REPLY)**

<https://blog.codecentric.de/en/2016/12/serverless-soap-legacy-api-integration-java-aws-lambda-aws-api-gateway/>

#### NEW QUESTION: 117

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

**Answer:** [\(SHOW ANSWER\)](#)

#### NEW QUESTION: 118

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

How can the developer accomplish this?

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

**Answer:** [B \(LEAVE A REPLY\)](#)

Reference:

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

#### NEW QUESTION: 119

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 handler. Place the connection in the global space.
- D. Replace RDS with Amazon DynamoDB to implement control over the number of writes per second.

**Answer: C (LEAVE A REPLY)**

Reference:

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: 120**

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

**Answer: A (LEAVE A REPLY)**

#### **NEW QUESTION: 121**

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

**Answer: (SHOW ANSWER)**

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (610 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### **NEW QUESTION: 122**

An existing serverless application processes uploaded image files. The process currently uses a single Lambda function that takes an image file, performs the processing, and stores the file in Amazon S3. Users of the application now require thumbnail generation of the images. Users want to avoid any impact to the time it takes to perform the image uploads.

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

**A.** Change the existing Lambda function handling the uploads to create thumbnails at the time of upload.

Have the function store both the image and thumbnail in Amazon S3.

**B.** Create an S3 event notification with a Lambda function destination. Create a new Lambda function to generate and store thumbnails.

**C.** Create an S3 event notification to an SQS Queue. Create a scheduled Lambda function that processes the queue, and generates and stores thumbnails.

**D.** Create a second Lambda function that handles thumbnail generation and storage. Change the existing Lambda function to invoke it asynchronously.

**Answer: A ([LEAVE A REPLY](#))**

### **NEW QUESTION: 123**

An application stores payroll information nightly in DynamoDB for a large number of employees across hundreds of offices. Item attributes consist of individual name, office identifier, and cumulative daily hours.

Managers run reports for ranges of names working in their office. One query is. "Return all Items in this office for names starting with A through E".

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

**A.** Configure the table to have a hash index on the name attribute, and a range index on the office identifier

**B.** Configure the table to have a range index on the name attribute, and a hash index on the office identifier

**C.** Configure a hash index on the name attribute and no range index

**D.** Configure a hash index on the office Identifier attribute and no range index

**Answer: ([SHOW ANSWER](#))**

Explanation

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.CoreComponents.html> Partition key and sort key - Referred to as a composite primary key, this type of key is composed of two attributes. The first attribute is the partition key, and the second attribute is the sort key. DynamoDB uses the partition key value as input to an internal hash function. The output from the hash function determines the partition (physical storage internal to DynamoDB) in which the item will be stored. All items with the same partition key value are stored together, in sorted order by sort key value.

### **NEW QUESTION: 124**

A company needs a version control system for collaborative software development. Features of the system must include the following:

\* Support for batches of changes across multiple files

\* Parallel branching

\* Version tracking

Which AWS service will meet these requirements?

**A.** AWS CodePipeline

**B.** Amazon S3

**C.** AWS Code Build

**D.** AWS CodeCommit

**Answer: ([SHOW ANSWER](#))**

<https://docs.aws.amazon.com/codecommit/latest/userguide/welcome.html>

#### NEW QUESTION: 125

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:
GsbAaIV7wlfj8zUqebHUANHzFbmkzILlxyj__y9xwhIhk99U_cUq1FieZnskWDjQ1wSHStVfdCEyZILGoccGpCiC
[hORceWF9rRwFTnEcRJ3N9iTrPAE1WHveC5Z54ALPaWlEjHlLg8CaB8d8lCKmxQuylCm0r1Bf2fHJRUjAYopMVMga
3olFmKA19yn_Z5rI120Q9p5ZIMX28zYM4dTulcJQUQjosgrEejfiIMYDda817Ooko9H6VmGJX62KfkRa5l7yE6hhf
?bIwA6tpyCJy2LWFRTe4bafqAyoqkarhPA4mGiZyWn4qSqb08e-
idSivWYPweaKGkampa0arcFR4gBD7Ph097WYBkzX9hViGddLMv41bXRv
```

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

- A. Use the AWS STS decode-authorization-message API to decode the message.
- B. Make a call to AWS KMS 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: C** ([LEAVE A REPLY](#))

#### NEW QUESTION: 126

An application uploads photos to an Amazon S3 bucket. Each photo that is uploaded to the S3 bucket must be resized to a thumbnail image by the application. Each thumbnail image is uploaded with a new name in the same S3 bucket. Which AWS service can a developer configure to directly process each single S3 event for each S3 object upload?

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

**Answer: B** ([LEAVE A REPLY](#))

#### NEW QUESTION: 127

A company is migrating a legacy application to a serverless application on AWS. The legacy application consists of a set of web services that are exposed by an Amazon API Gateway API. A developer needs to replace the existing implementation of web services with AWS Lambda functions. The developer needs to test a new version of the API that uses the functions in production. The developer must minimize the impact of the testing on the application's users.

Which solution will meet these requirements?

- A. Create a deployment stage. Enable mutual TLS for the new version of the API.
- B. Create a development stage for the new version of the API. Promote a canary release.
- C. Create a development stage for the new version of the API. Use a canary deployment.
- D. Create a beta stage for the new version of the API. Send the updated endpoint to the users.

**Answer: B** ([LEAVE A REPLY](#))

#### NEW QUESTION: 128

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

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

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

**Answer: C (LEAVE A REPLY)**

#### **NEW QUESTION: 129**

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

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

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

**Answer: (SHOW ANSWER)**

Reference:

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

#### **NEW QUESTION: 130**

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

What solution would give the BEST performance increase?

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

**Answer: A (LEAVE A REPLY)**

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

#### **NEW QUESTION: 131**

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

How can the developer accomplish this?

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

Answer: ([SHOW ANSWER](#))

#### NEW QUESTION: 132

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

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

The Developer needs to create/delete branches.

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

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

Answer: A ([LEAVE A REPLY](#))

#### NEW QUESTION: 133

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

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

**Answer: A** ([LEAVE A REPLY](#))

#### NEW QUESTION: 134

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": "21269319989900637946712965403778482371",
      "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. Reduce the number of KCL consumers.
- B. Use a PutRecord API instead of PutRecords
- C. Implement retries with exponential backoff
- D. Use Amazon SNS instead of Kinesis.
- E. Reduce the frequency and/or size of the requests

**Answer: C,E** ([LEAVE A REPLY](#))

#### NEW QUESTION: 135

A global company has an application running on Amazon EC2 instances that serves image files from Amazon S3. User requests from the browser are causing high traffic, which results in degraded performance.

Which optimization solution should a Developer implement to increase application performance?

- A. Create an Amazon ElastiCache cluster to cache and serve frequently accessed items.
- B. Create multiple prefix in the S3 bucket to increase the request rate
- C. Submit a ticket to AWS support to request a rate limit increase for the S3 bucket.
- D. Use Amazon CloudFront to serve the content of images stored in Amazon S3.

**Answer: D** ([LEAVE A REPLY](#))

**NEW QUESTION: 136**

Which DynamoDB limits can be raised by contacting AWS support? Choose 2 answers

- A. The number of hash keys per account
- B. The maximum storage used per account
- C. The number of tables per account
- D. The number of local secondary indexes per account
- E. The number of provisioned throughput units per account

**Answer: C,E** ([LEAVE A REPLY](#))

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

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (610 Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

**NEW QUESTION: 137**

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

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

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 138**

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

How can the Developer address this issue?

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

**Answer: C** ([LEAVE A REPLY](#))

**NEW QUESTION: 139**

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

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

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

**Answer: C (LEAVE A REPLY)**

#### **NEW QUESTION: 140**

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

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

**Answer: C,D (LEAVE A REPLY)**

#### **NEW QUESTION: 141**

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: (SHOW ANSWER)**

Reference:

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

#### **NEW QUESTION: 142**

A Developer created a Lambda function for a web application backend. When testing the Lambda function from the AWS Lambda console, the Developer can see that the function is being executed, but there is no log data being generated in Amazon CloudWatch Logs, even after several minutes.

What could cause this situation?

- A. The Lambda function does not have any explicit log statements for the log data to send it to CloudWatch Logs.
- B. The Lambda function is missing CloudWatch Logs as a source trigger to send log data.
- C. The execution role for the Lambda function is missing permissions to write log data to the CloudWatch Logs.
- D. The Lambda function is missing a target CloudWatch Log group.

**Answer: C (LEAVE A REPLY)**

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

### NEW QUESTION: 143

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

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

**Answer: A (LEAVE A REPLY)**

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

### NEW QUESTION: 144

A company is building an application to track athlete performance using an Amazon DynamoDB table. Each item in the table is identified by a partition key (user\_id) and a sort key (sport\_name). The table design is shown below:



(Note: Not all table attributes are shown)

A Developer is asked to write a leaderboard application to display the top performers (user\_id) based on the score for each sport\_name.

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

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

**Answer: B (LEAVE A REPLY)**

Explanation

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

[https://docs.aws.amazon.com/zh\\_cn/amazondynamodb/latest/developerguide/GSI.html](https://docs.aws.amazon.com/zh_cn/amazondynamodb/latest/developerguide/GSI.html)

### NEW QUESTION: 145

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

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

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 146**

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

**Answer: A** ([LEAVE A REPLY](#))

**NEW QUESTION: 147**

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

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

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

**Answer: (SHOW ANSWER)**

Reference:

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

**NEW QUESTION: 148**

A company needs a version control system for collaborative software development. Features of the system must include the following:

Support for batches of changes across multiple files

Parallel branching

Which AWS service will meet these requirements?

- A. AWS CodePipeline
- B. Amazon S3
- C. AWS Code Build
- D. AWS CodeCommit

**Answer: D** ([LEAVE A REPLY](#))

Reference:

<https://docs.aws.amazon.com/codecommit/latest/userguide/welcome.html>

**NEW QUESTION: 149**

A company needs a new REST API that can return information about the contents of an Amazon S3 bucket, such as a count of the objects stored in it.

The company has decided that the new API should be written as a microservice using AWS Lambda and Amazon API Gateway.

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

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

D. Create an IAM role, attach the AmazonS3FullAccess managed policy to it, and assign the role to the Lambda function as its execution role.

**Answer: B (LEAVE A REPLY)**

Explanation

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

#### NEW QUESTION: 150

While developing an application that runs on Amazon EC2 in an Amazon VPC, a Developer identifies the need for centralized storage of application-level logs.

Which AWS service can be used to securely store these logs?

- A. Amazon CloudWatch Logs
- B. AWS CloudTrail
- C. Amazon EC2 VPC Flow Logs
- D. Amazon CloudSearch

**Answer: A (LEAVE A REPLY)**

#### NEW QUESTION: 151

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

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

**Answer: C (LEAVE A REPLY)**

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (610 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 152

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

**Answer: C,E (LEAVE A REPLY)**

**NEW QUESTION: 153**

A company recently experienced some unexpected downtime. After investigating, the company determines that a developer mistakenly terminated several production Amazon EC2 instances. What should the company do to BEST protect against accidental terminations in the future.

- A. Enable EC2 termination protection on all production instances unless approval has been given through AWS Resource Access Manager.
- B. Enable EC2 termination protection on production instances. Deny the developer group's permissions policy access to terminate instance. Create a new role that developer can assume when termination is necessary.
- C. Modify the developer group's permissions policy to deny them access to delete production instances unless approved has been given through AWS Resource Access Manager.
- D. Modify the developer group's permission policy to require multi-factor authentication (MFA) only production instances are being delete Enable EC2 termination protection on production instances.

**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 154**

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

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

**Answer: A (LEAVE A REPLY)**

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

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

**NEW QUESTION: 155**

A developer has created a Node js web application on a local development machine. The developer wants to use AWS technology to host the website. The developer needs a solution that requires the least possible operational overhead and no code changes.

Which AWS service should the developer use to meet these requirements?

- A. AWS Elastic Beanstalk
- B. Amazon Elastic Kubernetes Service (Amazon EKS)
- C. AWS Lambda
- D. Amazon EC2

**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 156**

A developer works in an environment with multiple AWS accounts that have AWS Lambda functions processing the same 100 KB payloads. The developer wants to centralize the point of origin of the payloads to one account and have all the Lambda functions be invoked whenever the initiating event occurs in the parent account.

How can the developer design the workflow in the MOST efficient way, so all the multi-account Lambda functions get invoked when the event occurs?

- A.** Create a Lambda function in the parent account and use cross-account IAM roles with the AWS Security Token Service (AWS STS) AssumeRole API call to make AWS Lambda invoke the API call to invoke all the cross-account Lambda functions.
- B.** Subscribe all the multi-account Lambda functions to an Amazon SNS topic and make a SNS Publish API call with the payload to the SNS topic.
- C.** Set up an Amazon SQS queue with the queue policy permitting the ReceiveMessage action for multi-account Lambda functions. Then send the payload to the SQS queue using the sqs:SendMessage permission and poll the queue using multi-account Lambda functions.
- D.** Use a worker on an Amazon EC2 instance to poll for the payload event. Invoke all Lambda functions using the Lambda Invoke API after using cross-account IAM roles with the AWS Security Token Service (AWS STS) AssumeRole API call.

**Answer: B (LEAVE A REPLY)**

<https://aws.amazon.com/es/blogs/compute/cross-account-integration-with-amazon-sns/>

#### **NEW QUESTION: 157**

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.** 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
- B.** 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
- C.** 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.
- 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: (SHOW ANSWER)**

#### **NEW QUESTION: 158**

A developer is creating an application that is based on an AWS Lambda function The function uses the AWS SDK to read product price data from an Amazon S3 bucket and to write user information to an Amazon Aurora DB instance The Lambda function runs often, up to a few times each minute To meet performance requirements, the developer must minimize the run duration of the Lambda function Which actions can help the developer increase the performance? (Select TWO )

- A.** Use environment variables to pass operational parameters to the function.
- B.** Use most-restrictive permissions when setting the IAM policies for the Lambda IAM role.
- C.** Split the code into different Lambda functions to keep the functions smaller.
- D.** Read the S3 product price data initially, and cache it locally in the /tmp directory.
- E.** Initialize SDK clients and database connections outside of the function handler

**Answer: (SHOW ANSWER)**

#### **NEW QUESTION: 159**

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: B,D (LEAVE A REPLY)**

Explanation

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

#### **NEW QUESTION: 160**

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

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

**Answer: C (LEAVE A REPLY)**

#### **NEW QUESTION: 161**

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

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

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

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

**Answer: (SHOW ANSWER)**

#### **NEW QUESTION: 162**

A company wants to migrate an imaging service to Amazon EC2 while following security best practices. The images are sourced and read from a non-public Amazon S3 bucket. What should a developer do to meet these requirements?

- A. Create an EC2 service role with read-only permissions for the S3 bucket. Attach the role to the EC2 instance.
- B. Create an S3 service role with read-only permissions for the S3 bucket. Attach the role to the EC2 instance.
- C. Create an IAM user with read-only permissions for the S3 bucket. Temporarily store the user credentials in the Amazon EBS volume of the EC2 instance.

D. Create an IAM user with read-only permissions for the S3 bucket. Temporarily store the user credentials in the user data of the EC2 instance.

**Answer:** ([SHOW ANSWER](#))

#### NEW QUESTION: 163

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

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

A. Save the secrets in Amazon S3 with AWS KMS server-side encryption, and use a signed URL to access them by using the IAM role from Amazon EC2 instances.

B. Use the instance metadata to store the secrets and to programmatically access the secrets from EC2 instances.

C. Use the Amazon DynamoDB client-side encryption library to save the secrets in DynamoDB and to programmatically access the secrets from EC2 instances.

D. Use AWS SSM Parameter Store to store the secrets and to programmatically access them by using the IAM role from EC2 instances.

**Answer:** D ([LEAVE A REPLY](#))

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

#### NEW QUESTION: 164

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

A. Life-cycle management of IAM credentials

B. Decommissioning storage devices

C. Security Group and ACL (Access Control List) settings

D. Encryption of EBS (Elastic Block Storage) volumes

E. Controlling physical access to compute resources

F. Patch management on the EC2 instance's operating system

**Answer:** A,C,D,F ([LEAVE A REPLY](#))

Explanation

Physical and Environmental Security

AWS's data centers are state of the art, utilizing innovative architectural and engineering approaches.

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

\* When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals.

\* AWS uses the techniques detailed in DoD 5220.22-M (National Industrial Security Program Operating Manual) or NIST 800-88 (Guidelines for Media Sanitization) to destroy data as part of the decommissioning process.

\* All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

Storage Decommissioning

#### NEW QUESTION: 165

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 (LEAVE A REPLY)**

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

### NEW QUESTION: 166

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: (SHOW ANSWER)**

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 are 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.

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (610 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

### NEW QUESTION: 167

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:** ([SHOW ANSWER](#))

Reference:

<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.Errors.Handling>

### **NEW QUESTION: 168**

A customer wants to deploy its source code on an AWS Elastic Beanstalk environment. The customer needs to perform deployment with minimal outage and should only use existing instances to retain application access log.

What deployment policy would satisfy these requirements?

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

**Answer:** B ([LEAVE A REPLY](#))

### **NEW QUESTION: 169**

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

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

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

**Answer:** D ([LEAVE A REPLY](#))

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-use-s3-too.html#bp-use-s3-too-large-values>

For example, consider the ProductCatalog table in the Creating Tables and Loading Data for Code Examples in DynamoDB section. Items in this table store information about item price, description, book authors, and dimensions for other products. If you wanted to store an image of each product that was too large to fit in an item, you could store the images in Amazon S3 instead of in DynamoDB.

### **NEW QUESTION: 170**

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. Mappings
- B. Outputs

C. Parameters

D. Resources

**Answer:** [\(SHOW ANSWER\)](#)

#### **NEW QUESTION: 171**

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. Amazon SQS

C. Amazon API Gateway

D. AWS X-Ray

**Answer:** [B \(LEAVE A REPLY\)](#)

#### **NEW QUESTION: 172**

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 \(LEAVE A REPLY\)](#)

Explanation

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

#### **NEW QUESTION: 173**

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

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

A. Create an IAM User that allows write access to the DynamoDB table.

B. Add an IAM Role to a running EC2 instance.

C. Add an IAM User to a running EC2 Instance.

D. Launch an EC2 Instance with the IAM Role included in the launch configuration.

E. Create an IAM Role that allows write access to the DynamoDB table.

F. Launch an EC2 Instance with the IAM User included in the launch configuration.

**Answer:** [B,E \(LEAVE A REPLY\)](#)

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

#### **NEW QUESTION: 174**

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. Specify a shard iterator using the ShardIterator parameter.
- C. Implement exponential backoff on the GetRecords call and the PutRecords call.
- D. Increase the delay between the GetRecords call and the PutRecords call.
- E. Increase the number of shards in the data stream

**Answer:** [\(SHOW ANSWER\)](#)

#### **NEW QUESTION: 175**

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

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

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

**Answer:** [D \(LEAVE A REPLY\)](#)

#### **NEW QUESTION: 176**

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

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

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

**Answer:** [B \(LEAVE A REPLY\)](#)

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

#### **NEW QUESTION: 177**

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 \(LEAVE A REPLY\)](#)

Explanation

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

#### **NEW QUESTION: 178**

An application is being developed to audit several AWS accounts. The application will run in Account A and must access AWS services in Accounts B and C.

What is the MOST secure way to allow the application to call AWS services in each audited account?

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

**Answer:** [\(SHOW ANSWER\)](#)

[https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account-with-roles.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html)

#### **NEW QUESTION: 179**

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

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

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

**Answer:** [D \(LEAVE A REPLY\)](#)

#### **NEW QUESTION: 180**

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

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

**Answer:** [D \(LEAVE A REPLY\)](#)

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

#### **NEW QUESTION: 181**

Given the source code for an AWS Lambda function in the local store.py containing a handler function called get\_store and the following AWS CloudFormation template:

```
Transform: AWS::Serverless-2016-10-31
Resources:
  StoreFunc:
    Type: AWS::Serverless::Function
    Properties:
      Handler: store.get_store
      Runtime: python3.6
```

What should be done to prepare the template so that it can be deployed using the AWS CLI command `aws cloudformation deploy`?

- A. Use `aws cloudformation compile` to base64 encode and embed the source file into a modified CloudFormation template.
- B. Use `aws cloudformation package` to upload the source code to an Amazon S3 bucket and produce a modified CloudFormation template.
- C. Use `aws lambda zip` to package the source file together with the CloudFormation template and deploy the resulting zip archive.
- D. Use `aws serverless create-package` to embed the source file directly into the existing CloudFormation template.

**Answer:** ([SHOW ANSWER](#))

Explanation

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

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (**610** Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 182

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

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

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

**Answer:** ([SHOW ANSWER](#))

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

#### NEW QUESTION: 183

A company is creating a REST service using an Amazon API Gateway with AWS Lambda integration. The service must run different versions for testing purposes.

What would be the BEST way to accomplish this?

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

**Answer: C (LEAVE A REPLY)**

#### **NEW QUESTION: 184**

A company is migrating its on-premises database to Amazon RDS for MySQL. The company has read-heavy workloads, and wants to make sure it refactors its code to achieve optimum read performance for its queries.

How can this objective be met?

- A. Add a connection string to use a read replica on an EC2 instance.
- B. Add database retries to effectively use RDS with vertical scaling
- C. Use RDS with multi-AZ deployment
- D. Add a connection string to use an RDS read replica for read queries

**Answer: D (LEAVE A REPLY)**

#### **NEW QUESTION: 185**

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 (LEAVE A REPLY)**

Reference:

Reference: [https://docs.aws.amazon.com/aws-technical-content/latest/microservices-on-aws/microservices-on-aws.pdf?](https://docs.aws.amazon.com/aws-technical-content/latest/microservices-on-aws/microservices-on-aws.pdf?icmpid=link_from_whitepapers_page)

icmpid=link\_from\_whitepapers\_page (14)

#### **NEW QUESTION: 186**

A developer is building a web and mobile application for two types of users regular users and guest users Regular users are required to log in, but guest users do not log in Users should see only their data, regardless of whether they authenticate Users need AWS credentials before they can access AWS resources What is the MOST secure solution that the developer can implement to allow access for guest users?

- A. Use an Amazon Cognito credentials provider to issue temporary credentials that are linked to an unauthenticated role that has access to the required resources
- B. Generate temporary credentials Store the temporary credentials in AWS Secrets Manager Use the temporary credentials to access the required resources
- C. Set up an IAM user that has permissions to the required resources. Hardcode the IAM credentials in the web and mobile application.
- D. Generate temporary keys that are stored in AWS Key Management Service (AWS KMS) Use the temporary keys to access the required resources

**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 187**

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

**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 188**

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 (LEAVE A REPLY)**

Explanation

<https://aws.amazon.com/premiumsupport/knowledge-center/throttled-ddb/>

**NEW QUESTION: 189**

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: (SHOW ANSWER)**

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: 190**

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

What solution would give the BEST performance increase?

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

**Answer: A (LEAVE A REPLY)**

Explanation

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

**NEW QUESTION: 191**

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

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 192**

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

**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 193**

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

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

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

D. Serialize the image and store it in multiple DynamoDB tables

**Answer: C (LEAVE A REPLY)**

#### NEW QUESTION: 194

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

**Answer: A,E (LEAVE A REPLY)**

#### NEW QUESTION: 195

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

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 196

When developing an AWS Lambda function that processes Amazon Kinesis Data Streams, Administrators within the company must receive a notice that includes the processed data.

How should the Developer write the function to send processed data to the Administrators?

- A. Separate the Lambda handler from the core logic
- B. Use Amazon CloudWatch Events to send the processed data
- C. Publish the processed data to an Amazon SNS topic
- D. Push the processed data to Amazon SQS

**Answer: (SHOW ANSWER)**

Explanation

<https://stackoverflow.com/questions/13681213/what-is-the-difference-between-amazon-sns-and-amazon-sqs>

<https://stackoverflow.com/questions/31484868/can-you-publish-a-message-to-an-sns-topic-using-an-aws-lambda>

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (**610 Q&As Dumps, 40%OFF Special Discount: Exam-Tests**)

#### **NEW QUESTION: 197**

The Developer for a retail company must integrate a fraud detection solution into the order processing solution. The fraud detection solution takes between ten and thirty minutes to verify an order. At peak, the web site can receive one hundred orders per minute.

What is the most scalable method to add the fraud detection solution to the order processing pipeline?

- A.** Add all new orders to an SQS queue. Configure an Auto Scaling group that uses the queue depth metric as its unit of scale to launch a dynamically-sized fleet of EC2 instances spanning multiple AZs with the fraud detection solution installed on them to pull orders from this queue. Update the order with a pass or fails status.
- B.** Add all new orders to an Amazon Kinesis Stream. Subscribe a Lambda function to automatically read batches of records from the Kinesis Stream. The Lambda function includes the fraud detection software and will update the order with a pass or fail status.
- C.** Write all new orders to Amazon DynamoDB. Configure DynamoDB Streams to include all new orders. Subscribe a Lambda function to automatically read batches of records from the Kinesis Stream. The Lambda function includes the fraud detection software and will update the order with a pass or fail status.
- D.** Add all new orders to an Amazon SQS queue. Configure a fleet of 10 EC2 instances spanning multiple AZs with the fraud detection solution installed on them to pull orders from this queue. Update the order with a pass or fails status.

**Answer: C (LEAVE A REPLY)**

#### **NEW QUESTION: 198**

A company is building a serverless application that uses AWS Lambda

a. The application includes Lambda functions that are exposed by Amazon API Gateway. The functions will use several large third-party libraries, and the build artifacts will exceed 50 MB in size.

Which combination of steps should a developer take to prepare and perform the deployment? (Select TWO.)

- A.** Issue the `aws lambda update-function-code` CLI command with the `-s3-bucket` and `-s3-key` parameters.
- B.** Upload the build artifact to Amazon S3.
- C.** Issue the `aws cloudformation package` CLI command.
- D.** Issue the `aws lambda update-function-code` CLI command with a parameter that points to the source code in AWS CodeCommit.
- E.** Issue the `aws lambda update-function-code` CLI command with the `-zip-file fileb://my-function.zip` parameter.

**Answer: A,B (LEAVE A REPLY)**

#### **NEW QUESTION: 199**

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

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

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

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 200**

Which of the following are valid arguments for an SNS Publish request? Choose 3 answers

- A.** TopicArn
- B.** Subject
- C.** Destination
- D.** Format
- E.** Message
- F.** Language

**Answer:** **A,B,E** ([LEAVE A REPLY](#))

Explanation

[https://docs.aws.amazon.com/sns/latest/api/API\\_Publish.html](https://docs.aws.amazon.com/sns/latest/api/API_Publish.html)

Message

MessageAttributes

MessageStructure

PhoneNumber

Subject

TargetArn

TopicArn

#### **NEW QUESTION: 201**

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

**Answer:** **D** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 202**

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 (LEAVE A REPLY)**

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

#### NEW QUESTION: 203

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 direct 10% of users to the newly deployed version. If too many errors are encountered, send 100% of traffic to the previous version
- B. Do not make any changes to the application Deploy the new version of the code. If too many errors are encountered, point the application back to the previous version using the version number in the Amazon Resource Name (ARN)
- C. 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. If too many errors are encountered, point the alias back to the previous version
- D. 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: C (LEAVE A REPLY)**

#### NEW QUESTION: 204

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

**Answer: C (LEAVE A REPLY)**

**NEW QUESTION: 205**

A Developer is building a mobile application and needs any update to user profile data to be pushed to all devices accessing the specific identity. The Developer does not want to manage a back end to maintain the user profile data.

What is the MOST efficient way for the Developer to achieve these requirements using Amazon Cognito?

- A. Use Cognito federated identities.
- B. Use a Cognito user pool.
- C. Use Cognito Sync.
- D. Use Cognito events.

**Answer: (SHOW ANSWER)**

Reference:

Amazon Cognito Sync is an AWS service and client library that enables cross-device syncing of application-related user data. You can use it to synchronize user profile data across mobile devices and the web without requiring your own backend.

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

**NEW QUESTION: 206**

A legacy service has an XML-based SOAP interface. The Developer wants to expose the functionality of the service to external clients with the Amazon API Gateway. Which technique will accomplish this?

- A. Create a RESTful API with the API Gateway; transform the incoming JSON into a valid XML message for the SOAP interface using mapping templates.
- B. Create a RESTful API with the API Gateway; pass the incoming JSON to the SOAP interface through an Application Load Balancer.
- C. Create a RESTful API with the API Gateway; pass the incoming XML to the SOAP interface through an Application Load Balancer.
- D. Create a RESTful API with the API Gateway; transform the incoming XML into a valid message for the SOAP interface using mapping templates.

**Answer: A (LEAVE A REPLY)**

<https://blog.codecentric.de/en/2016/12/serverless-soap-legacy-api-integration-java-aws-lambda-aws-api-gateway/>

**NEW QUESTION: 207**

An application needs to encrypt data that is written to Amazon S3 where the keys are managed in an on-premises data center and the encryption is handled by S3. Which type of encryption should be used?

- A. Use server-side encryption with customer-provided keys
- B. Use client-side encryption with customer master keys
- C. Use server-side encryption with Amazon S3-managed keys
- D. Use server-side encryption with AWS KMS-managed keys.

**Answer: C (LEAVE A REPLY)**

**NEW QUESTION: 208**

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 (LEAVE A REPLY)**

Reference:

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

#### NEW QUESTION: 209

A company is building an application to track athlete performance using an Amazon DynamoDB table. Each item in the table is identified by a partition key (user\_id) and a sort key (sport\_name). The table design is shown below:



(Note: Not all table attributes are shown)

A Developer is asked to write a leaderboard application to display the top performers (user\_id) based on the score for each sport\_name.

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

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

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 210

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

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

**Answer: (SHOW ANSWER)**

#### NEW QUESTION: 211

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

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

- A. Install SSL certificates on the EC2 instances.
- B. Configure SSL certificates on an Elastic Load Balancer.

- C. Configure an Elastic Load Balancer with SSL termination.
- D. Configure an Elastic Load Balancer with SSL pass-through.
- E. Configure an Elastic Load Balancer with a Loadable Storage System.

**Answer:** ([SHOW ANSWER](#))

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (**610 Q&As Dumps, 40%OFF** **Special Discount: Exam-Tests**)

#### **NEW QUESTION: 212**

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

**Answer: D** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 213**

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** ([LEAVE A REPLY](#))

Reference:

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

#### **NEW QUESTION: 214**

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. A Docker in an application package
- C. An Amazon ECS cluster

D. A CLI for Elastic Beanstalk

Answer: A ([LEAVE A REPLY](#))

**Valid DVA-C01 Dumps** shared by TrainingQuiz.com for Helping Passing DVA-C01 Exam! TrainingQuiz.com now offer the **newest DVA-C01 exam dumps**, the TrainingQuiz.com DVA-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com DVA-C01 dumps with Test Engine here: <https://www.trainingquiz.com/DVA-C01-practice-quiz.html> (**610** Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)