

Exam Questions DVA-C01

AWS Certified Developer Associate Exam

<https://www.2passeasy.com/dumps/DVA-C01/>



NEW QUESTION 1

- (Exam Topic 4)

A company is building a serverless application that uses AWS Lambda. 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 `-zip-file fileb://my-function.zip` parameter
- 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 the `-s3-bucket` and `-s3-key` parameters.
- E. Issue the `aws lambda update-function-code` CLI command with a parameter that points to the source code in AWS CodeCommit.

Answer: BD

NEW QUESTION 2

- (Exam Topic 4)

A company must encrypt sensitive data that the company will store in Amazon S3. A developer must retain total control over the company's AWS Key Management Service (AWS KMS) key and the company's data keys. The company currently uses an on-premises hardware security module (HSM) solution. The company wants to move its key management onto AWS.

Which solution will meet these requirements?

- A. Implement server-side encryption with AWS KMS managed keys (SSE-KMS). Use AWS CloudHSM to generate the KMS key and data keys to use with AWS KMS.
- B. Implement server-side encryption with customer-provided encryption keys (SSE-C). Use AWS CloudHSM to generate the KMS key and manage the data keys that the company will use to read and write objects to Amazon S3.
- C. Implement server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use AWS CloudHSM to generate the KMS key and manage the data keys that the company will use to read and write objects to Amazon S3.
- D. Implement server-side encryption with AWS KMS managed keys (SSE-KMS). Use the AWS KMS custom key store feature to manage the data key
- E. Then read or write objects to Amazon S3 as normal.

Answer: D

Explanation:

<https://docs.aws.amazon.com/cloudhsm/latest/userguide/best-practices.html> Q: Can other AWS services use CloudHSM to store and manage keys? AWS services integrate with AWS Key Management Service, which in turn is integrated with AWS CloudHSM through the KMS custom key store feature. If you want to use the server-side encryption offered by many AWS services (such as EBS, S3, or Amazon RDS), you can do so by configuring a custom key store in AWS KMS.

NEW QUESTION 3

- (Exam Topic 4)

A developer is working on a Python application that runs on Amazon EC2 instances. The developer wants to enable tracing of application requests to debug performance issues in the code.

Which combination of actions should the developer take to achieve this goal? (Select TWO.)

- A. Install the Amazon CloudWatch agent on the EC2 instances.
- B. Install the AWS X-Ray daemon on the EC2 instances
- C. Configure the application to write JSON-formatted logs to `/var/log/cloudwatch`.
- D. Configure the application to write trace data to `/var/log/xray`.
- E. Install and configure the AWS X-Ray SDK for Python in the application.

Answer: BE

NEW QUESTION 4

- (Exam Topic 4)

A company is building an application for stock trading. The application needs sub-millisecond latency for processing trade requests. The company uses Amazon DynamoDB to store all the trading data that is used to process each trading request

A development team performs load testing on the application and finds that the data retrieval time is higher than expected. The development team needs a solution that reduces the data retrieval time with the least possible effort.

Which solution meets these requirements?

- A. Add local secondary indexes (LSIs) for the trading data
- B. Store the trading data in Amazon S3, and use S3 Transfer Acceleration.
- C. Add retries with exponential backoff for DynamoDB queries.
- D. Use DynamoDB Accelerator (DAX) to cache the trading data

Answer: D

NEW QUESTION 5

- (Exam Topic 4)

A developer is creating a serverless application that uses an AWS Lambda function. The developer will use AWS CloudFormation to deploy the application. The application will write logs to Amazon CloudWatch Logs. The developer has created a log group in a CloudFormation template for the application to use. The developer needs to modify the CloudFormation template to make the name of the log group available to the application at runtime.

Which solution will meet this requirement?

- A. Use the `AWS::Include` transform in CloudFormation to provide the log group's name to the application.
- B. Pass the log group's name to the application in the user data section of the CloudFormation template.

- C. Use the CloudFormation template's Mappings section to specify the log group's name for the application.
- D. Pass the log group's Amazon Resource Name (ARN) as an environment variable to the Lambda function.

Answer: C

NEW QUESTION 6

- (Exam Topic 4)

A developer has an application that is composed of many different AWS Lambda functions. The Lambda functions all use some of the same dependencies. To avoid security issues, the developer is constantly updating the dependencies of all of the Lambda functions. The result is duplicated effort for each function. Now can the developer keep the dependencies of the Lambda functions up to date with the LEAST additional complexity?

- A. Define a maintenance window for the Lambda functions to ensure that the functions get updated copies of the dependencies.
- B. Upgrade the Lambda functions to the most recent runtime version.
- C. Define a Lambda layer that contains all of the shared dependencies.
- D. Use an AWS CodeCommit repository to host the dependencies in a centralized location.

Answer: C

NEW QUESTION 7

- (Exam Topic 4)

A media company wants to test its web application more frequently. The company deploys the application by using a separate AWS CloudFormation stack for each environment. The same CloudFormation template is deployed to each stack as the application progresses through the development lifecycle. A developer needs to build an automated alert for the quality assurance (QA) team. The developer wants the alert to occur for new deployments in the final pre-production environment. Which solution will moot these requirements?

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic
- B. Add a subscription to notify the QA team
- C. Update the CloudFormation stack options to point to the SNS topic in the pro-production environment
- D. Most Voted
- E. Create an AWS Lambda function that notifies the QA team
- F. Create an Amazon EventBridge rule to invoke the Lambda function on the default event bus
- G. Filter the events on the CloudFormation service and the CloudFormation stack Amazon Resource Name (ARN).
- H. Create an Amazon CloudWatch alarm that monitors the metrics from CloudFormation
- I. Filter the metrics on the stack name and the stack status
- J. Configure the alarm to notify the QA team.
- K. Create an AWS Lambda function that notifies the QA team
- L. Configure the event source mapping to receive events from CloudFormation
- M. Specify the filtering values to limit invocations to the desired CloudFormation stack.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-rollback-email/>

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-rollback-email/> <https://www.trendmicro.com/cloudoneconformity/knowledge-base/aws/CloudFormation/cloudformation-stack-n>

NEW QUESTION 8

- (Exam Topic 4)

Which solution will meet these requirements?

- A. Build the container from the amazon/aws-xray-daemon base image
- B. Use the AWS X-Ray SDK to instrument the application.
- C. Install the Amazon CloudWatch agent on the container image
- D. Use the CloudWatch SDK to publish custom metrics from each of the microservices.
- E. Install the AWS X-Ray daemon on each of the ECS instances.
- F. Configure AWS CloudTrail data events to capture the traffic between the microservices.

Answer: C

NEW QUESTION 9

- (Exam Topic 4)

A developer is migrating a Windows-based legacy application from on premises to AWS. The application will run on Amazon EC2 instances that run Amazon Linux. The application stores a large number of files in an NFS drive. The migration solution must minimize downtime and application code changes. Which solution should the developer use to migrate the application data?

- A. Create an Amazon S3 bucket
- B. Use the s3 sync command to upload the files to the S3 bucket.
- C. Create an Amazon Elastic Block Store (Amazon EBS) volume
- D. Upload the files to the volume
- E. Attach the volume to the EC2 instances.
- F. Create an Amazon Elastic File System (Amazon EFS) file system
- G. Use AWS DataSync to transfer the files to Amazon EFS.
- H. Create an Amazon Elastic File System (Amazon EFS) file system
- I. Mount the EFS file system from the legacy application
- J. Copy the files to the EFS mount.

Answer: C

NEW QUESTION 10

- (Exam Topic 4)

A developer has an application container, an AWS Lambda function, and an Amazon Simple Queue Service (Amazon SQS) queue. The Lambda function uses the SQS queue as an event source. The Lambda function makes a call to a third-party machine learning API when the function is invoked. The response from the third-party API can take up to 60 seconds to return.

The Lambda function's timeout value is currently 65 seconds. The developer has noticed that the Lambda function sometimes processes duplicate messages from the SQS queue.

What should the developer do to ensure that the Lambda function does not process duplicate messages?

- A. Configure the Lambda function with a larger amount of memory.
- B. Configure an increase in the Lambda function's timeout value.
- C. Configure the SQS queue's delivery delay value to be greater than the maximum time it takes to call the third-party API.
- D. Configure the SQS queue's visibility timeout value to be greater than the maximum time it takes to call the third-party API.

Answer: A

NEW QUESTION 10

- (Exam Topic 4)

A developer designed an application on an Amazon EC2 instance. The application makes API requests to objects in an Amazon S3 bucket. Which combination of steps will ensure that the application makes the API requests in the MOST secure manner? (Select TWO.)

- A. Create an IAM user that has permissions to the S3 bucket.
- B. Add the user to an IAM group.
- C. Create an IAM role that has permissions to the S3 bucket.
- D. Add the IAM role to an instance profile.
- E. Attach the instance profile to the EC2 instance.
- F. Create an IAM role that has permissions to the S3 bucket.
- G. Assign the role to an IAM group.
- H. Store the credentials of the IAM user in the environment variables on the EC2 instance.

Answer: BC

NEW QUESTION 13

- (Exam Topic 4)

An application that is running on Amazon EC2 instances stores data in an Amazon S3 bucket. All the data must be encrypted in transit. How can a developer ensure that all traffic to the S3 bucket is encrypted?

- A. Install certificates on the EC2 instances.
- B. Create a private VPC endpoint.
- C. Configure the S3 bucket with server-side encryption with AWS KMS managed encryption keys (SSE-KMS).
- D. Create an S3 bucket policy that denies traffic when the value for the `aws:SecureTransport` condition key is false.

Answer: C

NEW QUESTION 14

- (Exam Topic 4)

A developer is monitoring an application that runs on an Amazon EC2 instance. The developer has configured a custom Amazon CloudWatch metric with a granularity of 1 second. If any issues occur, the developer wants to be notified within 30 seconds by Amazon Simple Notification Service (Amazon SNS).

What should the developer do to meet this requirement?

- A. Configure a high-resolution CloudWatch alarm.
- B. Set up a custom CloudWatch dashboard.
- C. Use Amazon CloudWatch Logs Insights.
- D. Change to a default CloudWatch metric.

Answer: D

NEW QUESTION 19

- (Exam Topic 4)

A company is using an Amazon API Gateway REST API endpoint as a webhook to publish events from an on-premises source control management (SCM) system to Amazon EventBridge. The company has configured an EventBridge rule to listen for the events and to control application deployment in a central AWS account. The company needs to receive the same events across multiple receiver AWS accounts.

How can a developer meet these requirements without changing the configuration of the SCM system?

- A. Deploy the API Gateway REST API to all the required AWS accounts.
- B. Use the same custom domain name for all the gateway endpoints so that a single SCM webhook can be used for all events from all accounts.
- C. Deploy the API Gateway REST API to all the receiver AWS accounts. Create as many SCM webhooks as the number of AWS accounts.
- D. Grant permission to the central AWS account for EventBridge to access the receiver AWS accounts. Add an EventBridge event bus on the receiver AWS accounts as the targets to the existing EventBridge rule.
- E. Convert the API Gateway type from REST API to HTTP API.

Answer: A

NEW QUESTION 20

- (Exam Topic 4)

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

Answer: B

NEW QUESTION 23

- (Exam Topic 4)

A developer is designing a serverless application that customers use to select seats for a concert venue. Customers send the ticket requests to an Amazon API Gateway API with an AWS Lambda function that acknowledges the order and generates an order ID. The application includes two additional Lambda functions: one for inventory management and one for payment processing. These two Lambda functions run in parallel and write the order to an Amazon Dynamo DB table. The application must provide seats to customers according to the following requirements. If a seat is accidentally sold more than once, the first order that the application received must get the seat. In these cases, the application must process the payment for only the first order. However, if the first order is rejected during payment processing, the second order must get the seat. In these cases, the application must process the payment for the second order.

Which solution will meet these requirements?

- A. Send the order ID to an Amazon Simple Notification Service (Amazon SNS) FIFO topic that fans out to one Amazon Simple Queue Service (Amazon SQS) FIFO queue for inventory management and another SQS FIFO queue for payment processing.
- B. Change the Lambda function that generates the order ID to initiate the Lambda function for inventory management
- C. Then initiate the Lambda function for payment processing.
- D. Send the order ID to an Amazon Simple Notification Service (Amazon SNS) topic
- E. Subscribe the Lambda functions for inventory management and payment processing to the topic.
- F. Deliver the order ID to an Amazon Simple Queue Service (Amazon SQS) queue
- G. Configure the Lambda functions for inventory management and payment processing to poll the queue.

Answer: A

Explanation:

Inventory & Payment functions are running in parallel. So going with Fanout option. <https://docs.aws.amazon.com/sns/latest/dg/sns-common-scenarios.html>

NEW QUESTION 28

- (Exam Topic 4)

A business intelligence application runs on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. Application-level audits require a searchable log of all API calls from users to the application. The application's developers must store the logs centrally on AWS.

Which solution will meet these requirements?

- A. Install the Amazon CloudWatch agent on the Amazon EC2 host that runs Fargate.
- B. Configure the awslogs log driver in the ECS task definition.
- C. Configure AWS CloudTrail for the ECS containers.
- D. Install the ECS logs collector on the ECS hosts.

Answer: B

Explanation:

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/using_awslogs.html

Configuring the awslogs log driver in the ECS task definition will allow the application to store the logs centrally on AWS. The awslogs log driver sends logs to Amazon CloudWatch Logs, which is a managed service that provides search and analysis of log data. This solution will meet the requirements of storing the logs centrally on AWS and making them searchable. Installing the Amazon CloudWatch agent on the Amazon EC2 host or installing the ECS logs collector on the ECS hosts will not work because the application is running on AWS Fargate and not on Amazon EC2. AWS CloudTrail is not a suitable solution because it is used to record API calls made to AWS services, not application-level API calls.

NEW QUESTION 33

- (Exam Topic 4)

A developer is creating an AWS CloudFormation template to deploy Amazon EC2 instances across multiple AWS accounts. The developer must choose the EC2 instances from a list of approved instance types.

How can the developer incorporate the list of approved instance types in the CloudFormation template?

- A. Create a separate CloudFormation template for each EC2 instance type in the list
- B. In the Resources section of the CloudFormation template, create resources for each EC2 instance type in the list.
- C. In the CloudFormation template, create a separate parameter for each EC2 instance type in the list.
- D. In the CloudFormation template, create a parameter with the list of EC2 instance types as AllowedValues

Answer: B

NEW QUESTION 36

- (Exam Topic 4)

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

Answer: C

Explanation:

This solution will meet the requirements most cost-effectively because it allows the developer to use a Lambda@Edge function to optimize the photos on the fly, without the need to pre-generate multiple variants of the photos for different devices. This approach can reduce the overall storage and compute costs associated with generating and storing multiple photo variants. Additionally, changing the CloudFront TTL cache policy to the maximum value possible can help reduce the number of times the Lambda@Edge function needs to be executed, further reducing the cost.
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Expiration.html>

NEW QUESTION 38

- (Exam Topic 4)

A company uses the AWS SDK for JavaScript in the Browser to build a web application and then hosts the application on Amazon S3. The company wants the application to support 10,000 users concurrently. The company selects Amazon DynamoDB to store user preferences in a table. There is a requirement to uniquely identify users at any scale.

Which solution will meet these requirements?

- A. Create a user cookie
- B. Attach an IAM role to the S3 bucket that hosts the application.
- C. Deploy an Amazon CloudFront distribution with an origin access identity (OAI) to access the S3 bucket.
- D. Configure and use Amazon Cognito
- E. Access DynamoDB with the authenticated users.
- F. Create an IAM user for each user
- G. Use fine-grained access control on the DynamoDB table to control access.

Answer: C

Explanation:

This will allow the application to support 10,000 users concurrently and will provide a unique identifier for each user. By using Amazon Cognito, the company can authenticate users and then access DynamoDB with the authenticated users to store their preferences in a table. This approach will allow the company to control access to the DynamoDB table and to scale to any number of users. Creating a user cookie or deploying an Amazon CloudFront distribution with an OAI would not solve the problem because these solutions do not provide a way to uniquely identify users or control access to DynamoDB. Creating an IAM user for each user and using fine-grained access control on the DynamoDB table would not be practical or scalable because it would require the company to manage and maintain a large number of IAM users.

When dealing with user profiles in serverless applications we often turn to Cognito for managing their credentials while the app itself will store user entities.

<https://www.sorenandersen.com/manage-user-profile-data-between-cognito-and-dynamodb/>

NEW QUESTION 41

- (Exam Topic 4)

A developer is working on an application that is deployed on an Amazon EC2 instance. The application needs to transfer a file to an Amazon S3 bucket. What should the developer do to authenticate the application's access to the S3 bucket in the MOST secure way?

- A. Create an access key for an IAM user
- B. Store the access key in the application's environment variables.
- C. Create an IAM role
- D. Create an access key for the role
- E. Store the access key in the application's environment variables.
- F. Associate an IAM role with the EC2 instance
- G. Use the instance metadata service to retrieve the credentials.
- H. Configure a bucket policy for the S3 bucket
- I. Allow access from the EC2 instance ID in the bucket policy.

Answer: B

NEW QUESTION 42

- (Exam Topic 4)

A developer deploys an ecommerce application on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group. The EC2 instances are based on an Amazon Machine Image (AMI) that uses an Amazon Elastic Block Store (Amazon EBS) root volume. After deployment, the developer notices that a third of the instances seem to be idle. These instances are not receiving requests from the load balancer. The developer verifies that all the instances are registered with the load balancer. The developer must implement a solution to allow the EC2 instances to receive requests from the load balancer.

Which action will meet this requirement?

- A. Reregister the failed instances with the ALB.
- B. Enable all Availability Zones for the ALB.
- C. Use the instance refresh feature to redeploy the EC2 Auto Scaling group.
- D. Restart the EC2 instances that are not receiving traffic.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/compute/introducing-instance-refresh-for-ec2-auto-scaling/>

NEW QUESTION 44

- (Exam Topic 4)

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

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 45

- (Exam Topic 4)

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

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

Answer: A

NEW QUESTION 47

- (Exam Topic 4)

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

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

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

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

Answer: B

NEW QUESTION 52

- (Exam Topic 4)

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

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

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

Answer: B

NEW QUESTION 54

- (Exam Topic 3)

A developer is building a backend system for the long-term storage of information from an inventory management system. The information needs to be stored so that other teams can build tools to report and analyze the data

How should the developer implement this solution to achieve the FASTEST running time?

- A. Create an AWS Lambda function that writes to Amazon S3 synchronously Increase the function's concurrency to match the highest expected value of concurrent scans and requests.
- B. Create an AWS Lambda function that writes to Amazon S3 asynchronously Configure a dead-letter queue to collect unsuccessful invocations
- C. Create an AWS Lambda function that writes to Amazon S3 synchronously Set the inventory system to retry failed requests.
- D. Create an AWS Lambda function that writes to an Amazon ElastiCache for Redis cluster asynchronously Configure a dead-letter queue to collect unsuccessful invocations.

Answer: A

NEW QUESTION 57

- (Exam Topic 3)

A company is using an AWS Lambda function to process records from an Amazon Kinesis data stream The company recently observed slow processing of the records. A developer notices that the iterator age metric for the function is Increasing and that the Lambda run duration is constantly above normal.

Which actions should the developer take to increase the processing speed? (Select TWO.)

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

Answer: DE

NEW QUESTION 62

- (Exam Topic 3)

A developer must increase read performance from an unencrypted Amazon S3 bucket. The application requires 100.000 read requests each second Cost-effectiveness is a priority. What would be the SIMPLEST approach to implement these requirements?

- A. Create 20 or more prefixes in Amazon S3 Place files by prefix
- B. Read in parallel by prefixes
- C. Create 20 or more AWS accounts Create a bucket in each account Read in parallel by bucket
- D. Deploy Memcached on Amazon EC2 Cache the files in memory Retrieve from the Memcached cache
- E. Copy all files to Amazon DynamoDB Index the files with S3 metadata Retrieve from DynamoDB

Answer: A

NEW QUESTION 63

- (Exam Topic 3)

A company has a large number of documents that are stored securely in Amazon S3 The company is creating an application that occasionally will read these documents The application will be deployed on Amazon EC2 instances.

The company's security requirements mandate that no long-term credentials can be stored on the EC2 instances and that only the needed documents can be accessed Only authorized users and applications can access the documents access must be logged by Amazon S3, and each document must follow S3 Lifecycle policies for archival and destruction

What should a developer do to meet these requirements?

- A. Create an event to invoke an AWS Lambda function when a document is uploaded Configure the function to write the documents to an Amazon Elastic File System (Amazon EFS) file system Configure the EC2 instances to mount the EFS file system Configure the application to access the documents that are stored in the file system as needed
- B. Create a user that has programmatic credentials, and attach a policy that allows read access to the S3 bucket Use the AWS CLI to configure those credentials for the EC2 instances to use Create an Amazon Machine Image (AMI), and add the access key and secret access key to the user data section to create environment variables Use the AMI to launch each EC2 instance that runs the application Add application code to use the keys that are stored in the environment variables to access the S3 bucket objects as needed.
- C. Modify the S3 bucket, make the bucket public, and make each object public Add application code to make REST calls to access the objects in the S3 bucket as needed
- D. Create an IAM role with permissions to read objects from Amazon S3 Attach the role to the EC2 instances as an instance profile Add application code to access the objects in the S3 bucket as needed.

Answer: D

NEW QUESTION 65

- (Exam Topic 3)

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

What is the MOST efficient method for performing this task?

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

Answer: D

NEW QUESTION 66

- (Exam Topic 3)

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

Which solution will meet this requirement?

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

Answer: C

NEW QUESTION 69

- (Exam Topic 3)

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

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

Which of the following solution supports these requirements?

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

Answer: C

NEW QUESTION 71

- (Exam Topic 3)

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. Create another Lambda subscription with the SNS message attribute value matching a filter option to 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 numeric value matching a filter option to save the appropriate orders to a separate table
- D. Modify the Lambda code to filter the orders and save the appropriate orders to a separate table

Answer: D

NEW QUESTION 74

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 76

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 81

- (Exam Topic 3)

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

- A. Deploy the application in a single Elastic Beanstalk environment

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

Answer: A

NEW QUESTION 86

- (Exam Topic 3)

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

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

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

Answer: C

NEW QUESTION 91

- (Exam Topic 3)

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

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

Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 95

- (Exam Topic 3)

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

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

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

Answer: DE

NEW QUESTION 99

- (Exam Topic 3)

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

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

What is the MOST likely cause of this issue?

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

Answer: B

NEW QUESTION 103

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 107

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 111

- (Exam Topic 3)

How does Envelope Encryption work in AWS KMS?

- A. The Customer Master Key is used to encrypt/decrypt a data key. The Plaintext Data Key is used to encrypt customer data.
- B. Two encryption keys are used. The Customer Master Key encrypts customer data.
- C. The Data Key is used to re-encrypt the encrypted data.
- D. Two encryption keys are used. The Data Key encrypts customer data. The Customer Master Key is used to re-encrypt the encrypted data.
- E. The Customer Master Key is used to encrypt/decrypt a data key.
- F. The Encrypted Data Key is used to encrypt customer data.

Answer: A

NEW QUESTION 114

- (Exam Topic 3)

A gaming application stores scores for players in an Amazon DynamoDB table that has four attributes: user_id, user_name, user_score, and user_rank. The users are allowed to update their names only. A user is authenticated by web identity federation.

Which set of conditions should be added in the policy attached to the role for the dynamodb:PutItem API call?

A)

```
"Condition": {
  "ForAllValues:StringEquals": {
    "dynamodb:LeadingKeys": [
      "${www.amazon.com:user_id}"
    ],
    "dynamodb:Attributes": [
      "user_name"
    ]
  }
}
```

B)

```
"Condition": {
  "ForAllValues:StringEquals": {
    "dynamodb:LeadingKeys": [
      "${www.amazon.com:user_name}"
    ],
    "dynamodb:Attributes": [
      "user_id"
    ]
  }
}
```

C)

```
"Condition": {
  "ForAllValues:StringEquals": {
    "dynamodb:LeadingKeys": [
      "${www.amazon.com:user_id}"
    ],
    "dynamodb:Attributes": [
      "user_name", "user_id"
    ]
  }
}
```

D)

```
"Condition": {
  "ForAllValues:StringEquals": {
    "dynamodb:LeadingKeys": [
      "${www.amazon.com:user_name}"
    ],
    "dynamodb:Attributes": [
      "user_name", "user_id"
    ]
  }
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

NEW QUESTION 116

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 120

- (Exam Topic 3)

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

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

Answer: BD

NEW QUESTION 122

- (Exam Topic 3)

A developer is designing a web application in which new users will use their email addresses to create accounts. Millions of users are expected to sign up. The application will store attributes for each user. Which AWS service or feature should the developer implement to meet these requirements?

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

Answer: A

NEW QUESTION 127

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 131

- (Exam Topic 3)

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

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

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

Answer: C

NEW QUESTION 132

- (Exam Topic 3)

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

What is the MOST operationally efficient way for the developer to improve the application's performance?

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

Answer: D

NEW QUESTION 136

- (Exam Topic 3)

A developer wants to insert a record into an Amazon DynamoDB table as soon as a new file is added to an Amazon S3 bucket.

Which set of steps would be necessary to achieve this?

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

Answer: B

NEW QUESTION 140

- (Exam Topic 3)

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

How can the developer improve the performance of these requests? (Select TWO)

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

Answer: AD

NEW QUESTION 141

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 144

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 146

- (Exam Topic 3)

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

Answer: AC

NEW QUESTION 147

- (Exam Topic 3)

A developer is building a highly secure healthcare application using .NET. The application requires writing temporary data to /tmp storage on an AWS Lambda function. How should the developer encrypt this data?

- A. Enable Amazon EBS volume encryption with an AWS KMS .NET configuration so that all storage attached to the Lambda function is encrypted.
- B. Set up the Lambda function with a role and key policy to access an AWS KMS CMK. Use the CMK to generate a data key used to encrypt all data prior to writing to /tmp storage.
- C. Use OpenSSL to generate a symmetric encryption key on Lambda startup. Use this key to encrypt the data prior to writing to /tmp.
- D. Use an on-premises hardware security module (HSM) to generate keys where the Lambda function requests a data key from the HSM and uses that to encrypt data on all requests to the function.

Answer: D

NEW QUESTION 149

- (Exam Topic 3)

A company is using continuous integration/continuous delivery (CI/CD) system. A developer must automate the deployment of an application software package to Amazon EC2 instances and virtual servers that run on premises. Which AWS services should the developer use to meet these requirements?

- A. AWS Cloud9
- B. AWS CodeBuild
- C. AWS Elastic Beanstalk
- D. AWS CodeDeploy

Answer: D

NEW QUESTION 152

- (Exam Topic 3)

A company processes incoming documents from an Amazon S3 bucket. Users upload documents to an S3 bucket using a web user interface. Upon receiving files in S3, an AWS Lambda function is invoked to process the files, but the Lambda function times out intermittently. If the Lambda function is configured with the default settings, what will happen to the S3 event when there is a timeout exception?

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

Answer: D

NEW QUESTION 154

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 158

- (Exam Topic 3)

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

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

Answer: AD

NEW QUESTION 162

- (Exam Topic 3)

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

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

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

Answer: AD

NEW QUESTION 163

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 167

- (Exam Topic 3)

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

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

What is the MOST likely cause of these issues?

- A. The command is incorrect; it should be rewritten to use : ut-i t am with a string argument.
- B. The developer needs to log a ticket with AWS Support to enable access to the demoman-table.
- C. Amazon DynamoDB cannot be accessed from the AWS CLI and needs to be called via the REST API
- D. The IAM user needs an associated policy with read access to demoman-table.

Answer: D

NEW QUESTION 172

- (Exam Topic 3)

A developer is working on a serverless application that needs to process any changes to an Amazon DynamoDB table with an AWS Lambda function

How should the developer configure the Lambda function to detect changes to the DynamoDB table?

- A. Create an Amazon Kinesis data stream, and attach it to the DynamoDB table. Create a trigger to connect the data stream to the Lambda function.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function on a regular schedule. Connect to the DynamoDB table from the Lambda function to detect changes.
- C. Enable DynamoDB Streams on the table. Create a trigger to connect the DynamoDB stream to the Lambda function.
- D. Create an Amazon Kinesis Data Firehose delivery stream, and attach it to the DynamoDB table. Configure the delivery stream destination as the Lambda function.

Answer: C

NEW QUESTION 176

- (Exam Topic 3)

A company is adding items to an Amazon DynamoDB table from an AWS Lambda function that is written in Python. A developer needs to implement a solution that inserts records in the DynamoDB table and performs automatic retry when the insert fails.

Which solution meets these requirements with MINIMUM code changes?

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

Answer: D

NEW QUESTION 179

- (Exam Topic 3)

An AWS Lambda function accesses two Amazon DynamoDB tables. A developer wants to improve the performance of the Lambda function by identifying bottlenecks in the function. How can the developer inspect the timing of the DynamoDB API calls?

- A. Add DynamoDB as an event source to the Lambda function.
- B. View the performance with Amazon CloudWatch metrics.
- C. Place an Application Load Balancer (ALB) in front of the two DynamoDB tables.

- D. Inspect the ALB logs
- E. Limit Lambda to no more than five concurrent invocations Monitor from the Lambda console
- F. Enable AWS X-Ray tracing for the functio
- G. View the traces from the X-Ray service.

Answer: D

NEW QUESTION 181

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 184

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 185

- (Exam Topic 3)

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

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

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

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

Answer: D

Explanation:

Meaning:

DOC-EXAMPLE-BUCKET ==> bucket

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

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

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

NEW QUESTION 189

- (Exam Topic 3)

A developer is working on a serverless project based in Java. Initial testing shows a cold start takes about 8 seconds on average for AWS Lambda functions. What should the developer do to reduce the cold start time? (Select TWO)

- A. Add the Spring Framework to the project and enable dependency injection
- B. Reduce the deployment package by including only the needed modules from the AWS SDK for Java.
- C. Increase the memory allocation setting for the Lambda function.

- D. Increase the timeout setting for the Lambda function.
- E. Change the Lambda invocation mode from synchronous to asynchronous.

Answer: BC

NEW QUESTION 192

- (Exam Topic 2)

A Developer is migrating an on-premises application to AWS. The application currently takes user uploads and saves them to a local directory on the server. All uploads must be saved and made immediately available to all instances in an Auto scaling group. Which approach will meet these requirements?

- A. Use Amazon EBS and configure the application AMI to use a snapshot of the same EBS instance on boot.
- B. Use Amazon S3 and rearchitect the application so all uploads are placed in S3.
- C. Use instance storage and share it between instances launched from the same Amazon machine image (AMI).
- D. Use Amazon EBS and file synchronization software to achieve eventual consistency among the auto scaling group.

Answer: B

Explanation:

Use Amazon S3 and rearchitect the application so all uploads are placed in S3. Even though you could do EBS attachment to ASG launch config userdata for ec2 instances going to serve, But you need to select the ASG in single AZ where your EBS is located otherwise it will not work since EBS is AZ locked.

NEW QUESTION 193

- (Exam Topic 2)

A gaming company is developing a mobile game application for iOS® and Android® platforms. This mobile game securely stores user data locally on the device. The company wants to allow users to use multiple device for the game, which requires user data synchronization across device. Which service should be used to synchronize user data across devices without the need to create a backend application?

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

Answer: D

NEW QUESTION 196

- (Exam Topic 2)

A Developer is building a web application that uses Amazon API Gateway to expose an AWS Lambda function to process requests from clients. During testing, the Developer notices that the API Gateway times out even though the Lambda function finishes under the set time limit. Which of the following API Gateway metrics in Amazon CloudWatch can help the Developer troubleshoot the issue? (Choose two.)

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

Answer: BC

Explanation:

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

NEW QUESTION 198

- (Exam Topic 2)

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

Answer: B

NEW QUESTION 200

- (Exam Topic 2)

A developer is storing sensitive data generated by an application in Amazon S3. The developer wants to encrypt the data at rest. A company policy requires an audit trail of when the master key was used and by whom. Which encryption option will meet these requirements?

- A. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- B. Server-side encryption with AWS KMS managed keys (SSE-KMS)
- C. Server-side encryption with customer-provided keys (SSE-C)
- D. Server-side encryption with self-managed keys

Answer: B

NEW QUESTION 204

- (Exam Topic 2)

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

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 208

- (Exam Topic 2)

A development team is creating a new application designed to run on AWS. While the test and production environments will run on Amazon EC2 instances, developers will each run their own environment on their laptops.

Which of the following is the simplest and MOST secure way to access AWS services from the local development machines?

- A. Use an IAM role to assume a role and execute API calls using the role.
- B. Create an IAM user to be shared with the entire development team, provide the development team with the access key.
- C. Create an IAM user for each developer on the team: provide each developer with a unique access key
- D. Set up a federation through an Amazon Cognito user pool.

Answer: A

NEW QUESTION 213

- (Exam Topic 2)

A company requires that AWS Lambda functions written by developers log errors so system administrators can more effectively troubleshoot issues. What should the developers implement to meet this need?

- A. Publish errors to a dedicated Amazon SQS queue
- B. Create an Amazon CloudWatch Events event to trigger based on certain Lambda events.
- C. Report errors through logging statements in Lambda function code.
- D. Set up an Amazon SNS topic that sends logging statements upon failure

Answer: B

NEW QUESTION 216

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 220

- (Exam Topic 2)

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 221

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 225

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 229

- (Exam Topic 2)

A company is using continuous integration and continuous delivery systems. A Developer now needs to automate a software package deployment to both Amazon EC2 instances and virtual servers running on-premises.
Which AWS service should be used to accomplish this?

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

Answer: D

NEW QUESTION 234

- (Exam Topic 2)

A company has implemented AWS CodePipeline to automate its release pipelines. The development team is writing an AWS Lambda function that will send notifications for state changes of each of the actions in the stages.
Which steps must be taken to associate the Lambda function with the event source?

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

Answer: B

NEW QUESTION 237

- (Exam Topic 2)

A software company needs to make sure user-uploaded documents are securely stored in Amazon S3. The documents must be encrypted at rest in Amazon S3. The company does not want to manage the security infrastructure in-house, but the company still needs extra protection to ensure it has control over its encryption keys due to industry regulations.
Which encryption strategy should a developer use to meet these requirements?

- A. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- B. Server-side encryption with customer-provided encryption keys (SSE-C)
- C. Server-side encryption with AWS KMS managed keys (SSE-KMS)
- D. Client-side encryption

Answer: D

NEW QUESTION 240

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 241

- (Exam Topic 2)

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: C

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>

NEW QUESTION 243

- (Exam Topic 2)

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

What is the MOST cost-effective solution?

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

Answer: C

Explanation:

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

NEW QUESTION 245

- (Exam Topic 2)

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

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

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

Answer: A

NEW QUESTION 249

- (Exam Topic 2)

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

Explanation:

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

NEW QUESTION 250

- (Exam Topic 2)

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

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

D. Pass the access key id and secret access key as parameters for each AWS CLI command.

Answer: C

NEW QUESTION 253

- (Exam Topic 2)

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

Which solution BEST meets these requirements?

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

Answer: C

NEW QUESTION 256

- (Exam Topic 2)

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

Which approach will meet these requirements?

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

Answer: D

Explanation:

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

NEW QUESTION 260

- (Exam Topic 2)

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

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

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

Answer: B

NEW QUESTION 264

- (Exam Topic 2)

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

What is the fastest way to process all the files?

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

Answer: C

NEW QUESTION 267

- (Exam Topic 2)

A developer is implementing authentication and authorization for an application. The developer needs to ensure that the user credentials are never exposed. Which approach should the developer take to meet this requirement?

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

Answer: C

NEW QUESTION 272

- (Exam Topic 2)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 276

- (Exam Topic 2)

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

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

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

Answer: B

NEW QUESTION 277

- (Exam Topic 2)

A company provides APIs as a service and commits to a service level agreement (SLA) with all its users. To comply with each SLA, what should the company do?

- A. Enable throttling limits for each method in Amazon API Gateway.
- B. Create a usage plan for each user and request API keys to access the APIs.
- C. Enable API rate limiting in Amazon Cognito for each user.
- D. Enable default throttling limits for each stage after deploying the APIs.

Answer: D

NEW QUESTION 279

- (Exam Topic 2)

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

Why is the table being throttled?

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

Answer: A

Explanation:

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

NEW QUESTION 282

- (Exam Topic 2)

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: A

Explanation:

https://aws.amazon.com/s3/faqs/?nc1=h_ls.

NEW QUESTION 286

- (Exam Topic 2)

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
- B. Use an event notification to invoke a Lambda function to publish data to CloudWatch.
- C. Publish the metric data to an Amazon Kinesis Stream using a PutRecord API call
- D. Subscribe a Lambda function that publishes data to CloudWatch.
- E. Use the CloudWatch PutMetricData API call to submit a custom metric to CloudWatch
- F. Provide the required credentials to enable the API call.
- G. Use the CloudWatch PutMetricData API call to submit a custom metric to CloudWatch
- H. Launch the EC2 instance with the required IAM role to enable the API call.

Answer: D

Explanation:

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

NEW QUESTION 288

- (Exam Topic 2)

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

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

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

Answer: CD

NEW QUESTION 290

- (Exam Topic 2)

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

What should the Development team do to achieve this?

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

Answer: A

Explanation:

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

NEW QUESTION 291

- (Exam Topic 2)

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

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

Answer: B

Explanation:

AWS credential profile, and enabling the Git credential helper to send the path to repositories: Reference:

<https://docs.aws.amazon.com/codecommit/latest/userguide/setting-up-https-unixes.html>

NEW QUESTION 294

- (Exam Topic 2)

An on-premises application is implemented using a Linux, Apache, MySQL and PHP (LAMP) stack. The Developer wants to run this application in AWS.

Which of the following sets of AWS services can be used to run this stack?

- A. Amazon API Gateway, Amazon S3
- B. AWS Lambda, Amazon DynamoDB
- C. Amazon EC2, Amazon Aurora
- D. Amazon Cognito, Amazon RDS
- E. Amazon ECS, Amazon EBS

Answer: C

NEW QUESTION 299

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 300

- (Exam Topic 2)

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

What is the best option to store the session state?

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

Answer: A

Explanation:

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

NEW QUESTION 301

- (Exam Topic 2)

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

Which type of encryption meets these requirements?

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

Answer: B

NEW QUESTION 304

- (Exam Topic 2)

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

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 305

- (Exam Topic 2)

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

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

- A. Increase the Amazon SQS queue polling timeout

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

Answer: A

NEW QUESTION 306

- (Exam Topic 2)

A company has an application where reading objects from Amazon S3 is based on the type of user. The user types are registered user and guest user. The company has 25,000 users and is growing. Information is pulled from an S3 bucket depending on the user type. Which approaches are recommended to provide access to both user types? (Select TWO.)

- A. Provide a different access key and secret access key in the application code for registered users and guest users to provide read access to the objects
- B. Use S3 bucket policies to restrict read access to specific IAM users
- C. Use Amazon Cognito to provide access using authenticated and unauthenticated roles
- D. Create a new IAM user for each user and grant read access.
- E. Use the AWS IAM service and let the application assume the different roles using the AWS Security Token Service (AWS STS) AssumeRole action depending on the type of user and provide read access to Amazon S3 using the assumed role

Answer: BC

NEW QUESTION 307

- (Exam Topic 2)

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

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

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

Answer: C

NEW QUESTION 311

- (Exam Topic 2)

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

Which approach would provide the BEST user experience?

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

Answer: C

NEW QUESTION 315

- (Exam Topic 1)

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

What is the reason for the duplicates?

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

Answer: C

Explanation:

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

NEW QUESTION 317

- (Exam Topic 1)

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

What action would minimize the Lambda compute time consumed?

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

Answer: D

NEW QUESTION 319

- (Exam Topic 1)

A meteorological system monitors 600 temperature gauges, obtaining temperature samples every minute and saving each sample to a DynamoDB table. Each sample involves writing 1K of data and the writes are evenly distributed over time.

How much write throughput is required for the target table?

- A. 1 write capacity unit
- B. 10 write capacity units
- C. 60 write capacity units
- D. 600 write capacity units
- E. 3600 write capacity units

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h>

NEW QUESTION 320

- (Exam Topic 1)

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

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

Answer: AE

Explanation:

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

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

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

NEW QUESTION 321

- (Exam Topic 1)

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

Explanation:

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

NEW QUESTION 323

- (Exam Topic 1)

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

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

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

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

```
    db.close()
```

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

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

Answer: C

Explanation:

Refer AWS documentation - Lambda Best Practices

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

NEW QUESTION 326

- (Exam Topic 1)

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

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

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

Answer: A

NEW QUESTION 330

- (Exam Topic 1)

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 re-factors its code to achieve optimum read performance for its queries.

How can this objective be met?

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

Answer: C

NEW QUESTION 333

- (Exam Topic 1)

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

Which AWS service should be used to refactor this architecture?

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

Answer: D

NEW QUESTION 338

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 341

- (Exam Topic 1)

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

How can compliance with this policy be ensured?

- A. Use AWS Lambda to send notifications to the security team if unencrypted objects are pun in the bucket.
- B. Configure an Amazon S3 bucket policy to prevent the upload of objects that do not contain the x-amzserver-side-encryption header.

- C. Create an Amazon CloudWatch event rule to verify that all objects stored in the Amazon S3 bucket are encrypted.
- D. Configure an Amazon S3 bucket policy to prevent the upload of objects that contain the x-amz-server-side-encryption header.

Answer: B

NEW QUESTION 346

- (Exam Topic 1)

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: CE

Explanation:

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

NEW QUESTION 347

- (Exam Topic 1)

A Developer wants access to make the log data of an application running on an EC2 instance available to systems administrators. Which of the following enables monitoring of this metric in Amazon CloudWatch?

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

Answer: D

NEW QUESTION 351

- (Exam Topic 1)

A company has an internet-facing application that uses Web Identity Federation to obtain a temporary credential from AWS Security Token Service (AWS STS). The app then uses the token to access AWS services.

Review the following response:

```
<AssumeRoleWithWebIdentityResponse xmlns="https://sts.amazonaws.com/doc/2011-06-08/">
  <AssumeRoleWithWebIdentityResult>
    <SubjectFromWebIdentityToken>amzn1.account.AF6RHO7KZUSXRVQJGKX6HB56KR2A</SubjectFromWebIdentityToken>
    <Audience>client.5498841531868486423.1543@apps.example.com</Audience>
    <AssumedRoleUser>
      <Arn>arn:aws:sts::123456789012:assumed-role/FederatedWebIdentityRole/app1</Arn>
      <AssumedRoleId>AROACLKWSQRAOEXAMPLE:app1</AssumedRoleId>
    </AssumedRoleUser>
    <Credentials>
      <SessionToken>AQoDYXdzEEB0a9ANXDXOXKXNolewxE5T1jQyp+IEXAMPLE</SessionToken>
      <SecretAccessKey>wJalrXUtnFEMI/KTMDENG/bPxrRfiCYzEXAMPLEKEY</SecretAccessKey>
      <Expiration>2014-10-24T23:00:23Z</Expiration>
      <AccessKeyId>ASgeIAIOSFODNN7EXAMPLE</AccessKeyId>
    </Credentials>
    <Provider>www.amazon.com</Provider>
  </AssumeRoleWithWebIdentityResult>
  <ResponseMetadata>
    <RequestId>ad4156e9-bce1-11e2-82e6-6b6efEXAMPLE</RequestId>
  </ResponseMetadata>
</AssumeRoleWithWebIdentityResponse>
```

Based on the response displayed what permissions are associated with the call from the application?

- A. Permissions associated with the role AROACLKWSQRAOEXAMPLE:app1
- B. Permissions associated with the default role used when the AWS service was built
- C. Permission associated with the IAM principal that owns the AccessKeyId ASgeIAIOSFODNN7EXAMPLE
- D. Permissions associated with the account that owns the AWS service

Answer: C

NEW QUESTION 356

- (Exam Topic 1)

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

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

- A. An Amazon SNS topic to be triggered when a concurrent user threshold is met

- B. An Amazon Cloudwatch Networkin metric
- C. Amazon CloudFront to leverage AWS Edge Locations
- D. A Custom Amazon CloudWatch metric for concurrent users.

Answer: D

NEW QUESTION 358

- (Exam Topic 1)

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

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

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

Answer: D

NEW QUESTION 362

- (Exam Topic 1)

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

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

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

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

Answer: AC

NEW QUESTION 367

- (Exam Topic 1)

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

What could be an explanation for this?

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 370

- (Exam Topic 1)

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

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

Answer: C

NEW QUESTION 374

- (Exam Topic 1)

What item operation allows the retrieval of multiple items from a DynamoDB table in a single API call?

- A. GetItem
- B. BatchGetItem
- C. GetMultipleItems
- D. GetItemRange

Answer: B

Explanation:

https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_BatchGetItem.html

NEW QUESTION 379

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 383

- (Exam Topic 1)

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

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

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

Answer: D

NEW QUESTION 388

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 393

- (Exam Topic 1)

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

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

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

Answer: C

NEW QUESTION 394

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 399

- (Exam Topic 1)

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

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

Answer: ACDF

Explanation:

Physical and Environmental Security

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

Storage Decommissioning

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

NEW QUESTION 401

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DVA-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DVA-C01 Product From:

<https://www.2passeasy.com/dumps/DVA-C01/>

Money Back Guarantee

DVA-C01 Practice Exam Features:

- * DVA-C01 Questions and Answers Updated Frequently
- * DVA-C01 Practice Questions Verified by Expert Senior Certified Staff
- * DVA-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DVA-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year