

Amazon-Web-Services

Exam Questions DVA-C01

AWS Certified Developer Associate Exam



NEW QUESTION 1

- (Exam Topic 4)

A developer needs to secure the static assets in a company's Amazon S3 bucket that is named DOC-EXAMPLE-BUCKET. The company has an Amazon CloudFront distribution that serves the S3 bucket's assets to the public. The developer has already created the origin access identity (OAI) and has associated the OAI with the distribution. The developer must write a bucket policy that allows only the CloudFront distribution to access the S3 bucket

Which policy will meet this requirement MOST securely?

A)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access Identity E12345OAI12"
  },
  "Action": [
    "s3:GetObject"
  ],
  "Resource": [
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
  ]
}
```

B)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access Identity E12345OAI12"
  },
  "Action": [
    "s3:GetObject",
    "s3:PutObject"
  ],
  "Resource": [
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
  ]
}
```

C)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access Identity E12345OAI12"
  },
  "Action": [
    "s3:GetObject"
  ],
  "Resource": [
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET"
  ]
}
```

D)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access Identity E12345OAI12"
  },
  "Action": [
    "s3:GetObject",
    "s3:PutObject"
  ],
  "Resource": [
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET"
  ]
}
```

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

Answer: C

NEW QUESTION 2

- (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 3

- (Exam Topic 4)

A developer is deploying a company's application to Amazon EC2 instances. The application generates gigabytes of data files each day. The files are rarely accessed, but the files must be available to the application's users within minutes of a request during the first year of storage. The company must retain the files for 7 years.

How can the developer implement the application to meet these requirements MOST cost-effectively?

- A. Store the files in an Amazon S3 bucket. Use the S3 Glacier Instant Retrieval storage class.
- B. Create an S3 Lifecycle policy to transition the files to the S3 Glacier Deep Archive storage class after 1 year.
- C. Store the files in an Amazon S3 bucket.
- D. Use the S3 Standard storage class.
- E. Create an S3 Lifecycle policy to transition the files to the S3 Glacier Flexible Retrieval storage class after 1 year.
- F. Store the files on an Amazon Elastic Block Store (Amazon EBS) volume.
- G. Use Amazon Data Lifecycle Manager (Amazon DLM) to create snapshots of the EBS volumes and to store those snapshots in Amazon S3.
- H. Store the files on an Amazon Elastic File System (Amazon EFS) mount.
- I. Configure EFS lifecycle management to transition the files to the EFS Standard-Infrequent Access (Standard-IA) storage class after 1 year.

Answer: A

Explanation:

Amazon S3 Glacier Instant Retrieval is an archive storage class that delivers the lowest-cost storage for long-lived data that is rarely accessed and requires retrieval in milliseconds. With S3 Glacier Instant Retrieval, you can save up to 68% on storage costs compared to using the S3 Standard-Infrequent Access (S3 Standard-IA) storage class, when your data is accessed once per quarter.

<https://aws.amazon.com/s3/storage-classes/glacier/instant-retrieval/>

NEW QUESTION 4

- (Exam Topic 4)

A developer at a company recently created a serverless application to process and show data from business reports. The application's user interface (UI) allows users to select and start processing the files. The UI displays a message when the result is available to view. The application uses AWS Step Functions with AWS Lambda functions to process the files. The developer used Amazon API Gateway and Lambda functions to create an API to support the UI.

The company's UI team reports that the request to process a file is often returning timeout errors because of the size or complexity of the files. The UI team wants the API to provide an immediate response so that the UI can display a message while the files are being processed. The backend process that is invoked by the API needs to send an email message when the report processing is complete.

What should the developer do to configure the API to meet these requirements?

- A. Change the API Gateway route to add an `X-Amz-Invocation-Type` header with a static value of 'Event' in the integration request.
- B. Deploy the API Gateway stage to apply the changes.
- C. Change the configuration of the Lambda function that implements the request to process a file. Configure the maximum age of the event so that the Lambda function will run asynchronously.
- D. Change the API Gateway timeout value to match the Lambda function timeout value.
- E. Deploy the API Gateway stage to apply the changes.
- F. Change the API Gateway route to add an `X-Amz-Target` header with a static value of 'Async' in the integration request.
- G. Deploy the API Gateway stage to apply the changes.

Answer: A

NEW QUESTION 5

- (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 6

- (Exam Topic 4)

A developer needs to use Amazon DynamoDB to store customer orders. The developer's company requires all customer data to be encrypted at rest with a key the company generates.

What should the developer do to meet these requirements?

- A. Create the DynamoDB table with encryption set to Non
- B. Code the application to use the key to decrypt the data when the application reads from the table
- C. Code the application to use the key to encrypt the data when the application writes to the table.
- D. Store the key by using AWS Key Management Service (AWS KMS). Choose an AWS KMS customer managed key during creation of the DynamoDB table
- E. Provide the Amazon Resource Name (ARN) of the AWS KMS key.
- F. Store the key by using AWS Key Management Service (AWS KMS). Create the DynamoDB table with default encryption
- G. Include the kms:Encrypt parameter with the Amazon Resource Name (ARN) of the AWS KMS key when using the DynamoDB software development kit (SDK).
- H. Store the key by using AWS Key Management Service (AWS KMS). Choose an AWS KMS AWS managed key during creation of the DynamoDB table Provide the Amazon Resource Name (ARN) of the AWS KMS key.

Answer: C

NEW QUESTION 7

- (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 8

- (Exam Topic 4)

A distributed application includes an AWS Lambda function that runs successfully in the DEV environment with 128 MB of memory assigned. The same function is failing in the TEST environment. The developer is monitoring the application using AWS X-Ray, but the Lambda function cannot be seen on the X-Ray service graph. The Lambda execution role has AWS X-Ray permissions

What is the MOST LIKELY cause for AWS X-Ray not showing any data for the Lambda function?

- A. The AWS SDK needs to be included in the AWS Lambda deployment package.
- B. VPC Flow Logs are not enabled for the application VPC.
- C. Active tracing needs to be enabled for the Lambda function
- D. The memory needs to be increased to 2 GB for the TEST environments.

Answer: C

Explanation:

<https://stackoverflow.com/questions/43728674/enabling-x-ray-support-in-aws-lambda> You need to check the "Enable Active Tracing" checkbox in the Lambda console.

NEW QUESTION 9

- (Exam Topic 4)

A developer is using AWS CodeDeploy to automate a company's application deployments to Amazon EC2. Which application specification file properties are required to ensure the software deployments do not fail?

(Select TWO.)

- A. The file must be a JSON-formatted file named `appspec.json`.
- B. The file must be a YAML-formatted file named `appspec.yml`
- C. The file must be stored in AWS CodeBuikJ and referenced from the application's source code.
- D. The file must be placed in the root of the directory structure of the application's source code.
- E. The file must be stored in Amazon S3 and referenced from the application's source code.

Answer: BD

NEW QUESTION 10

- (Exam Topic 4)

A company is running an application on AWS Elastic Beanstalk in a single-instance environment. The company's deployments must avoid any downtime. Which deployment option will meet these requirements?

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

Answer: D

NEW QUESTION 10

- (Exam Topic 4)

A developer has built an application running on AWS Lambda using AWS Serverless Application Model (AWS SAM). What is the correct sequence of steps 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 CodeCommit.

Answer: B

NEW QUESTION 11

- (Exam Topic 4)

A developer is creating a serverless application that uses an AWS Lambda function. The developer will use AWS Cloud Formation 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 12

- (Exam Topic 4)

An open-source map application gathers data from several geolocation APIs. The application's source code repository is public and can be used by anyone, but the geolocation APIs must not be directly accessible.

A developer must implement a solution to prevent the credentials that are used to access the APIs from becoming public. The solution also must ensure that the application still functions properly.

Which solution will meet these requirements MOST cost-effectively?

- A. Store the credentials in AWS Secrets Manager
- B. Retrieve the credentials by using the GetSecretValue API operation.
- C. Store the credentials in AWS Key Management Service (AWS KMS). Retrieve the credentials by using the GetPublicKey API operation.
- D. Store the credentials in AWS Security Token Service (AWS STS). Retrieve the credentials by using the GetCallerIdentity API operation.
- E. Store the credentials in AWS Systems Manager Parameter Store
- F. Retrieve the credentials by using the GetParameter API operation.

Answer: D

Explanation:

Secrets Manager: It is paid. The storage cost is \$0.40 per secret per month and API interactions cost is \$0.05 per 10,000 API calls. Parameter Store: For Standard parameters, No additional charge for storage and standard throughput. For higher throughput, API interactions cost is \$0.05 per 10,000 API calls. For Advanced parameters, storage cost is \$0.05 per advanced parameter per month and API interactions cost is \$0.05 per 10,000 API calls.
<https://aws.amazon.com/systems-manager/pricing/>

NEW QUESTION 15

- (Exam Topic 4)

A company hosts a monolithic application on Amazon EC2 instances. The company starts converting some features of the application to a serverless architecture by using Amazon API Gateway and AWS Lambda. After the migration, some users report problems with payment processing.

Upon inspection, a developer discovers that the Lambda function that calls the external payment API is taking longer than expected. Therefore, the API Gateway requests are timing out.

What should the developer do to resolve this issue in the serverless architecture?

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

Answer: B

NEW QUESTION 16

- (Exam Topic 4)

A developer is receiving an intermittent ProvisionedThroughputExceededException error from an application that is based on Amazon DynamoDB. According to the Amazon CloudWatch metrics for the table, the application is not exceeding the provisioned throughput.

What could be the cause of the issue?

- A. The DynamoDB table storage size is larger than the provisioned size.
- B. The application is exceeding capacity on a particular hash key.
- C. The DynamoDB table is exceeding the provisioned scaling operations.
- D. The application is exceeding capacity on a particular sort key.

Answer: B

NEW QUESTION 17

- (Exam Topic 4)

A company requires objects that are stored in Amazon S3 to be encrypted. The company is currently using server-side encryption with AWS KMS managed encryption keys (SSE-KMS). A developer needs to optimize the cost-effectiveness of the encryption mechanism without negatively affecting performance. What should the developer do to meet these requirements?

- A. Change the encryption type to customer-provided keys.
- B. Configure the S3 bucket to use an S3 Bucket Key for SSE-KMS.
- C. Use S3 bucket policies to limit the principals who can create objects.
- D. Use a custom policy to limit the number of AWS KMS calls that are allowed.

Answer: B

NEW QUESTION 20

- (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 24

- (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 meet 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 pre-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 26

- (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 30

- (Exam Topic 4)

A developer notices timeouts from the AWS CLI when the developer runs list commands. What should the developer do to avoid these timeouts?

- A. Use the `--page-size` parameter to request a smaller number of items.
- B. Use shorthand syntax to separate the list by a single space.
- C. Use the `yaml-stream` output for faster viewing of large datasets.
- D. Use quotation marks around strings to enclose data structure.

Answer: A

NEW QUESTION 34

- (Exam Topic 4)

A company has an application that uses Amazon Cognito user pools as an identity provider. The company must secure access to user records. The company implements multi-factor authentication (MFA). The company also wants to send a login activity notification by email every time a user logs in. What is the MOST operationally efficient solution that meets this requirement?

- A. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification.
- B. Add an Amazon API Gateway to invoke the function.
- C. Call the API from the client side when login confirmation is received.
- D. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification.
- E. Add an Amazon Cognito post authentication Lambda trigger for the function.
- F. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification.
- G. Create an Amazon CloudWatch Logs log subscription filter to invoke the function based on the login status.
- H. Configure Amazon Cognito to stream all logs to Amazon Kinesis Data Firehose.
- I. Create an AWS Lambda function to process the streamed logs and it send the email notification based on the login status of each user.

Answer: C

NEW QUESTION 35

- (Exam Topic 4)

A developer is creating a serverless orchestrator that performs a series of steps to process incoming IoT data. The orchestrator transforms the data, performs a series of calculations, and stores the results in Amazon DynamoDB. The entire process is completed in less than a minute. The orchestrator must process tens of thousands of transactions each second. The orchestrator must not miss a transaction and must process each transaction at least once.

Which solution will meet these requirements MOST cost-effectively?

- A. Use Amazon Simple Notification Service (Amazon SNS) to process the data through an HTTP or HTTPS endpoint.
- B. Use AWS Step Functions to process the data as Standard Workflows.
- C. Use AWS Step Functions to process the data as Synchronous Express Workflows.
- D. Use AWS Step Functions to process the data as Asynchronous Express Workflows.

Answer: D

NEW QUESTION 38

- (Exam Topic 4)

A company uses Amazon DynamoDB as a data store for its order management system. The company's frontend application stores orders in a DynamoDB table. The DynamoDB table is configured to send change events to a DynamoDB stream. The company uses an AWS Lambda function to log and process the incoming orders based on data from the DynamoDB stream.

An operational review reveals that the order quantity of incoming orders is sometimes set to 0. A developer needs to create a dashboard that will show how many unique customers this problem affects each day.

What should the developer do to implement the dashboard?

- A. Grant the Lambda function's execution role permissions to upload logs to Amazon CloudWatch Logs. Implement a CloudWatch Logs Insights query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1-day period.
- B. Add the CloudWatch Logs Insights query to a CloudWatch dashboard.
- C. Use Amazon Athena to query AWS CloudTrail API logs for API call.
- D. Implement an Athena query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1-day period.
- E. Add the Athena query to an Amazon CloudWatch dashboard.
- F. Configure the Lambda function to send events to Amazon EventBridge.
- G. Create an EventBridge rule that groups the number of unique customers for orders with order quantity equal to 0 in 1-day period.
- H. Add a CloudWatch dashboard as the target of the rule.
- I. Turn on custom Amazon CloudWatch metrics for the DynamoDB stream of the DynamoDB table. Create a CloudWatch alarm that groups the number of unique customers for orders with order quantity equal to 0 in 1-day period.
- J. Add the CloudWatch alarm to a CloudWatch dashboard.

Answer: D

NEW QUESTION 42

- (Exam Topic 4)

A company that manages movie reviews wants to make its movie review data available to its customers by calling a set of REST web service endpoints. The company will develop the retrieval functionality as AWS Lambda functions and will expose the functionality to customers as an Amazon API Gateway REST API. The company needs to ensure that no consumer exceeds 100 requests a day to the API during the initial deployment. The company decides to use API Gateway API keys to restrict access. The company creates and issues API keys for each customer.

What should the company do next to meet these requirements with the LEAST administrative effort?

- A. Create a usage plan that applies throttling at 100 requests a day.
- B. Associate the usage plan with the API keys of all customers.
- C. Create an Amazon DynamoDB table to track all the requests that use a particular API key.
- D. For each request to the API, count the number of records in the DynamoDB table for that day for the API key.
- E. If the number of requests is 100 or greater, generate an exception.
- F. Create a usage plan that applies a quota of 100 requests a day.
- G. Associate the usage plan with the API keys of all customers.
- H. Create an Amazon Aurora table to track all the requests that use a particular API key.
- I. For each request to the API, count the number of records in the Aurora table for that day for the API key.
- J. If the number of requests is 100 or greater, generate an exception.

Answer: C

Explanation:

What is quota in API gateway?

API quotas usually describe a certain amount of calls for longer intervals. For example, your API quota might be 5,000 calls per month. Remember that this could be combined with a rate limit or throttling setup e.g. 20 TPS (Transactions per Second).

What is throttling in AWS API gateway?

AWS throttling limits are applied across all accounts and clients in a region. These limit settings exist to prevent your API—and your account—from being overwhelmed by too many requests. These limits are set by AWS and can't be changed by a customer.

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-create-usage-plans-with-console.html> Choose Enable quota, and set specify a value (for example, 5000) for a selected time interval (for example, Month).

NEW QUESTION 44

- (Exam Topic 4)

A developer has an application that pushes files from an on-premises local server to an Amazon S3 bucket. The application uses an AWS access key and a secret key that are stored on the server for authentication. The application calls AWS Security Token Service (AWS STS) to assume a role with access to perform the S3 PUT operation to upload the file.

The developer is migrating the server to an Amazon EC2 instance. The EC2 instance is configured with an IAM instance profile in the same AWS account that owns the S3 bucket.

What is the MOST secure solution for the developer to use to migrate the automation code?

- A. Remove the code that calls the STS AssumeRole operation
- B. Use the same access key and secret key from the server to access the S3 bucket.
- C. Remove the access key and the secret key
- D. Use the STS AssumeRole operation to add permissions to access the S3 bucket.
- E. Remove the access key, the secret key, and the code that calls the STS AssumeRole operation
- F. Use an IAM instance profile role that grants access to the S3 bucket.
- G. Remove the access key, the secret key, and the code that calls the STS AssumeRole operation
- H. Create a new access key and secret key
- I. Use the new keys to access the S3 bucket.

Answer: C

NEW QUESTION 45

- (Exam Topic 4)

A developer is writing an application that stores data in an Amazon DynamoDB table by using the PutItem API operation. The table has a partition key of streamID and has a sort key of seqID. The developer needs to make sure that the PutItem invocation does not overwrite the existing partition key and sort key.

Which condition expression will maintain the uniqueness of the partition key and the sort key?

A)

```
condition = 'attribute_not_exists(streamID)
AND
attribute_not_exists(seqID)'
```

B)

```
condition = 'attribute_not_exists(PARTITION)
AND
attribute_not_exists(SORT)'
```

C)

```
condition = 'attribute_exists(streamID)
AND
attribute_exists(seqID)'
```

D)

```
condition = 'attribute_exists(PARTITION)
AND
attribute_exists(SORT)'
```

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

Answer: A

Explanation:

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

NEW QUESTION 50

- (Exam Topic 4)

A developer wants to implement authentication using Amazon Cognito user pools for an existing API in Amazon API Gateway. After creating the Amazon Cognito user pool, the developer tests the GET request to the API. Unauthenticated requests to the API return a 200 OK status response.

Which combination of additional steps are required to complete the authentication implementation? (Select TWO.)

- A. Create an Amazon Cognito authorizer in API Gateway and specify the Amazon Cognito user pool.
- B. Create an AWS Lambda authorizer in API Gateway and specify the Amazon Cognito user pool.
- C. Specify the authorizer in the GET method section of API Gateway and redeploy the API.
- D. Use Amazon Cognito user pools to make and authenticate the request to API Gateway.

E. Create an Amazon Cognito authorizer in API Gateway and specify the Amazon Cognito identity pool.

Answer: AD

NEW QUESTION 55

- (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 60

- (Exam Topic 4)

A company manages a financial services application that stores a large volume of data in an Amazon DynamoDB table. A developer needs to improve the performance of the DynamoDB read queries without increasing the cost. Which solution meets these requirements?

- A. Use parallel scans
- B. Add a local secondary index (LSI).
- C. Create a DynamoDB Accelerator (DAX) cluster.
- D. Query with the Projection Expression parameter

Answer: C

NEW QUESTION 65

- (Exam Topic 4)

A company that has multiple offices uses an Amazon DynamoDB table to store employee payroll information. Item attributes consist of employee names, office identifiers, and cumulative daily hours worked. The most frequently used query extracts a report of an alphabetical subset of employees for a specific office. Which design of the DynamoDB table primary key will have the MINIMUM performance impact?

- A. Partition key on the office identifier and sort key on the employee name
- B. Partition key on the employee name and sort key on the office identifier
- C. Partition key on the employee name
- D. Partition key on the office identifier

Answer: B

NEW QUESTION 66

- (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 69

- (Exam Topic 4)

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 73

- (Exam Topic 4)

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

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

Answer: A

NEW QUESTION 74

- (Exam Topic 4)

A company has a front-end application that runs on four Amazon EC2 instances behind an Elastic Load Balancer (ELB) in a production environment that is provisioned by AWS Elastic Beanstalk. A developer needs to deploy and test new application code while updating the Elastic Beanstalk platform from the current version to a newer version of Node.js. The solution must result in zero downtime for the application.

Which solution meets these requirements?

- A. Clone the production environment to a different platform version
- B. Deploy the new application code, and test it
- C. Swap the environment URLs upon verification.
- D. Deploy the new application code in an all-at-once deployment to the existing EC2 instance
- E. Test the code
- F. Redeploy the previous code if verification fails.
- G. Perform an immutable update to deploy the new application code to new EC2 instance
- H. Serve traffic to the new instances after they pass health checks
- I. Use a rolling deployment for the new application code
- J. Apply the code to a subset of EC2 instances until the tests pass
- K. Redeploy the previous code if the tests fail.

Answer: D

NEW QUESTION 79

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

Answer: B

NEW QUESTION 84

- (Exam Topic 4)

A developer is creating a photo website. Amazon Route 53 hosts the website's domain. The developer wants to store the application code and images in an Amazon S3 bucket. The developer also wants to use Amazon CloudFront to deliver the images to users.

The developer has created the S3 bucket and a CloudFront distribution. The developer wants the images to be accessed only through the website's domain. Users must not use the S3 URLs.

Which solution will meet these requirements?

- A. Create a CloudFront origin access identity (OAI). Associate the OAI with the CloudFront distribution. Modify the S3 bucket policy to allow access from only the OAI
- B. Create an alias in Route 53 that points the website domain to the S3 bucket.
- C. Create a CloudFront origin access identity (OAI). Associate the OAI with the CloudFront distribution. Modify the S3 bucket policy to allow access from only the OAI
- D. Update the Route 53 records to point the website domain to the CloudFront domain name.
- E. Block public access in the S3 bucket policy
- F. Configure CloudFront to use the S3 bucket endpoint. Create an alias in Route 53 that points the website domain to the S3 bucket.
- G. Block public access in the S3 bucket policy
- H. Configure CloudFront to use the S3 bucket endpoint. Create an alias in Route 53 that points the website domain to the CloudFront domain name.

Answer: A

NEW QUESTION 88

- (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 90

- (Exam Topic 4)

An application adds a processing date to each transaction that it receives. The application writes each transaction to an Amazon DynamoDB table by using the PutItem operation. Each transaction has a unique ID (transactionID). Sometimes the application receives transactions more than once. A developer notices that duplicate transactions in DynamoDB have the latest processing date instead of the date when the transaction was first received. Duplicate records happen infrequently, and most of the transactions are unique. What is the MOST cost-effective solution that the developer can implement to ensure that PutItem does not update an existing record?

- A. Call the GetItem operation first to confirm that the record does not exist.
- B. Then call PutItem.
- C. Enable the TTL attribute on the DynamoDB table.
- D. Implement a conditional put by using the attribute_exists(transactionID) condition expression.
- E. Implement a conditional put by using the attribute_not_exists(transactionID) condition expression.

Answer: C

NEW QUESTION 92

- (Exam Topic 4)

A developer is creating a web application that collects highly regulated and confidential user data through a POST request. The web application is served through Amazon CloudFront. User names and phone numbers must be encrypted at the edge and must remain encrypted throughout the entire application stack. What is the MOST secure way to meet these requirements?

- A. Enforce Match Viewer with HTTPS Only on CloudFront.
- B. Use only the newest TLS security policy on CloudFront.
- C. Enforce a signed URL on CloudFront on the front end.
- D. Use field-level encryption on CloudFront.

Answer: D

Explanation:

Field-level encryption allows you to enable your users to securely upload sensitive information to your web servers. The sensitive information provided by your users is encrypted at the edge, close to the user, and remains encrypted throughout your entire application stack. <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/field-level-encryption.html>

NEW QUESTION 94

- (Exam Topic 4)

A company is running its website on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group. A developer needs to secure the internet-facing connection with HTTPS. The developer uses AWS Certificate Manager (ACM) to issue an X.509 certificate. What should the developer do to secure the connection?

- A. Configure the ALB to use the X.509 certificate by using the AWS Management Console.
- B. Configure each EC2 instance to use the same X.509 certificate by using the AWS Management Console.
- C. Export the root key of the X.509 certificate to an Amazon S3 bucket.
- D. Configure each EC2 instance to use the same X.509 certificate from the S3 bucket.
- E. Export the root key of the X.509 certificate to an Amazon S3 bucket.
- F. Configure the ALB to use the X.509 certificate from the S3 bucket.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/configure-acm-certificates-ec2/> <https://aws.amazon.com/premiumsupport/knowledge-center/associate-acm-certificate-alb-nlb/>
Configuring an Amazon Issued ACM public certificate for a website that's hosted on an EC2 instance requires exporting the certificate. However, you can't export the certificate because ACM manages the private key that signs and creates the certificate. Instead, you can associate an ACM certificate with a load balancer or an ACM SSL/TLS certificate with a CloudFront distribution.

Associate an ACM SSL certificate with an Application Load Balancer

Open the Amazon EC2 console.

In the navigation pane, choose Load Balancers, and then choose your Application Load Balancer. Choose Add listener.

For Protocol, choose HTTPS. For port, choose 443.

For Default action(s), choose Forward to, and then select your ALB target group from the dropdown list. For Default SSL certificate, choose From ACM (recommended) and then choose the ACM certificate. Choose Save.

NEW QUESTION 98

- (Exam Topic 4)

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

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

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

Which solution will meet these requirements?

- A. Add a second cache behavior to the distribution with the same origin as the default cache behavior
- B. Set the path pattern for the second cache behavior to the path of the login page, and make viewer access unrestricted
- C. Keep the default cache behavior's settings unchanged.
- D. Add a second cache behavior to the distribution with the same origin as the default cache behavior
- E. Set the path pattern for the second cache behavior to *, and make viewer access restricted
- F. Change the default cache behavior's path pattern to the path of the login page, and make viewer access unrestricted.
- G. Add a second origin as a failover origin to the default cache behavior
- H. Point the failover origin to the S3 bucket
- I. Set the path pattern for the primary origin to * and make viewer access restricted
- J. Set the path pattern for the failover origin to the path of the login page, and make viewer access unrestricted.
- K. Add a bucket policy to the S3 bucket to allow read access
- L. Set the resource on the policy to the Amazon Resource Name (ARN) of the login page object in the S3 bucket
- M. Add a CloudFront function to the default cache behavior to redirect unauthorized requests to the login page's S3 URI.

Answer: B

Explanation:

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

NEW QUESTION 103

- (Exam Topic 4)

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

- A. Create an AWS Lambda rotation function that has appropriate IAM permissions. Store the password in AWS Secrets Manager. Configure Secrets Manager to rotate the password by using the Lambda function.
- B. Encrypt the existing password with AWS Key Management Service (AWS KMS). Export the existing password. Generate a random password with AWS KMS. Use the AWS KMS password renewal feature to replace the existing password with the new password.
- C. Create an AWS Lambda rotation function that has appropriate IAM permissions. Store the password in AWS Systems Manager Parameter Store. Configure Parameter Store to rotate the password by using the Lambda function.
- D. Integrate AWS Systems Manager Parameter Store with a Key Management Interoperability Protocol (KMIP)-compliant third-party secret manager to enable third-party database password rotation on AWS.

Answer: C

NEW QUESTION 107

- (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 112

- (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 use
- 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 114

- (Exam Topic 4)

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

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

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

Answer: C

NEW QUESTION 117

- (Exam Topic 4)

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

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

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

Answer: AE

Explanation:

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

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

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

E) DynamoDB is also common to store session state with TTL support. And all user data stored in Amazon DynamoDB is fully encrypted at rest -

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

NEW QUESTION 122

- (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 124

- (Exam Topic 4)

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

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

Answer: A

NEW QUESTION 127

- (Exam Topic 4)

A developer is building varKHJS microservices for an application that will run on Amazon EC2 instances. The developer needs to monitor the end-to-end view of the requests between the microservices and debug any issues in the various microservices.

What should the developer do to accomplish these tasks?

- A. Use Amazon CloudWatch to aggregate the microservices' logs and metrics, and build the monitoring dashboard
- B. Use AWS CloudTrail to aggregate the microservices' logs and metrics, and build the monitoring dashboard
- C. Use the AWS X-Ray SDK to add instrumentation in all the microservices, and monitor using the X-Ray service map.
- D. Use AWS Health to monitor the health of all the microservices

Answer: C

Explanation:

<https://docs.aws.amazon.com/whitepapers/latest/microservices-on-aws/distributed-monitoring.html#distributed-t> <https://aws.amazon.com/xray/>

NEW QUESTION 128

- (Exam Topic 4)

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

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

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

Answer: B

NEW QUESTION 129

- (Exam Topic 4)

A developer is writing an application that will run on Amazon EC2 instances in an Auto Scaling group. The developer wants to externalize the session state to support the application. Which AWS services or resources can the developer use to meet these requirements? (Select TWO.)

- A. Amazon DynamoDB
- B. Amazon Cognito
- C. Amazon ElastiCache
- D. Application Load Balancer
- E. Amazon Simple Queue Service (Amazon SQS)

Answer: AC

NEW QUESTION 134

- (Exam Topic 4)

A company has a serverless application that uses AWS Lambda functions and AWS Systems Manager parameters to store configuration data. The company..... the Lambda functions inside the VPC and into private subnets. The Lambda functions are now producing errors in their attempts to access Systems Manager parameters.

Which solution will allow the Lambda functions to access Systems Manager parameters inside the VPC?

- A. Configure security groups to allow access to Systems Manager.
- B. Create an interface VPC endpoint for Systems Manager.
- C. Use an internet gateway from inside the VPC.
- D. Create a gateway VPC endpoint for Systems Manager.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/lambda-vpc-parameter-store/>

NEW QUESTION 135

- (Exam Topic 4)

A company is concerned that a malicious user could deploy unauthorized changes to the code for an AWS Lambda function. What can a developer do to ensure that only trusted code is deployed to Lambda?

- A. Turn on the trusted code option in AWS CodeDeploy
- B. Add the CodeDeploy digital certificate to the Lambda package before deploying the package to Lambda
- C. Define the code signing configuration in the Lambda console Use AWS Signer to digitally sign the Lambda package before deploying the package to Lambda
- D. Link Lambda to AWS Key Management Service (AWS KMS) in the Lambda console
- E. Use AWS KMS to digitally sign the Lambda package before deploying the package to Lambda.
- F. Set the KmsKeyArn property of the Lambda function to the Amazon Resource Name (ARN) of a trusted key before deploying the package to Lambda.

Answer: B

NEW QUESTION 137

- (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 140

- (Exam Topic 3)

A developer has an AWS CodePipeline pipeline that invokes AWS CodeBuild in the build stage The developer wants to pass in a variable from CodePipeline so that the variable can be read in the CodeBuild buildspec.yml file

How can the developer accomplish this goal?

- A. Configure a unique CodePipeline variable namespace and variables as key-value pairs that define each of the variables required in CodeBuild
- B. Configure a CodePipeline environment variable that contains a JSON document that defines each of the variables required in CodeBuild
- C. Configure an AWS CloudFormation stack set that contains a JSON document that defines each of the variables required in CodeBuild Reference the stack set from CodePipeline
- D. Configure an AWS CodeArtifact repository to store each environment variable Reference CodeArtifact from CodePipeline and CodeBuild

Answer: B

NEW QUESTION 141

- (Exam Topic 3)

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

Which approach will improve performance and MINIMIZE management overhead?

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

Answer: B

NEW QUESTION 146

- (Exam Topic 3)

A developer has written an application that runs on Amazon EC2 instances. The developer is adding

functionality for the application to write objects to an Amazon S3 bucket Which policy must the developer modify to allow the instances to write these objects?

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

Answer: A

NEW QUESTION 147

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 152

- (Exam Topic 3)

A developer from AnyCompany's AWS account needs access to the Example Corp AWS account AnyCompany uses an identity provider that is compatible with OpenID Connect.

What is the MOST secure way for Example Corp to allow developer access?

- A. Create a cross-account role and call the AssumeRole API operation
- B. Create a user in the Example Corp account and provide the access keys
- C. Create a user in the Example Corp account and provide the credentials
- D. Create a cross-account role and call the AssumeRoleWithWebIdentity API operation

Answer: B

NEW QUESTION 153

- (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 158

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 163

- (Exam Topic 3)

A company runs its APIs using Amazon API Gateway in front of AWS Lambda functions The company wants to add logging at the API level Each API must have production and development environments The developer wants to enable different logging levels in both environments.

How can these requirements be met?

- A. Set up a stage for each environment In each stage, point to different Lambda functions that implement the logging logic in the code Access the logs in Amazon CloudWatch Logs
- B. Set up a stage for each environment In each stage, define a different logging level according to the logging requirements Access the logs in Amazon CloudWatch Logs
- C. Set up a stage and use the same Lambda functions In Amazon CloudWatch Logs set up a filter based on the log level according to the logging requirements
- D. Set up a stage for each environment In each stage, define a variable for the log level Set the value according to the logging requirements.

Answer: A

NEW QUESTION 168

- (Exam Topic 3)

An application development team decides to use AWS X Ray to monitor application code to analyze performance and perform root cause analysis

What does the team need to do to begin using X Ray? (Select TWO)

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

Answer: DE

NEW QUESTION 169

- (Exam Topic 3)

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

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

- A. AWS Elastic Beanstalk
- B. Amazon EC2
- C. AWS Lambda

D. Amazon Elastic Kubernetes Service (Amazon EKS)

Answer: A

NEW QUESTION 173

- (Exam Topic 3)

A developer must modify an Alexa skill backed by an AWS Lambda function to access an Amazon DynamoDB table in a second account. A role in the second account has been created with permissions to access the table. How should the table be accessed?

- A. Modify the Lambda function execution role's permissions to include the new role
- B. Change the Lambda function execution role to be the new role
- C. Assume the new role in the Lambda function when accessing the table
- D. Store the access key and the secret key for the new role and use them when accessing the table

Answer: A

NEW QUESTION 178

- (Exam Topic 3)

A company is using AWS CloudFormation templates to deploy AWS resources. The company needs to update one of its AWS CloudFormation stacks. What can the company do to find out how the changes will impact the resources that are running?

- A. Investigate the change sets
- B. Investigate the stack policies
- C. Investigate the Metadata section.
- D. Investigate the Resources section.

Answer: A

NEW QUESTION 182

- (Exam Topic 3)

A developer is building an application on Amazon EC2. The developer encountered an "Access Denied" error on some of the API calls to AWS services while testing. The developer needs to modify permissions that have been already given to the instance. How can these requirements be met with minimal changes and minimum downtime?

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

Answer: D

NEW QUESTION 183

- (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 186

- (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 188

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 191

- (Exam Topic 3)

An application running on multiple Amazon EC2 instances pulls messages from an 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 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 196

- (Exam Topic 3)

A large company has its application components distributed across multiple AWS accounts. The company needs to collect and visualize trace data across these accounts.

What should be used to meet these requirements?

- A. AWS X-Ray
- B. Amazon CloudWatch
- C. Amazon VPC flow logs
- D. Amazon Elasticsearch Service

Answer: A

NEW QUESTION 199

- (Exam Topic 3)

A developer is building a static, client-side rendered website that is powered by ReactJS. The code has no server-side generated components and does not need to run any programming languages on the server. However, the code serves static HTML, CSS, and JavaScript to the client on each request. The developer's solution to host the website must maximize performance and cost-effectiveness.

Which combination of AWS services or resources should the developer use to meet these requirements?

- A. Application Load Balancer and Amazon EC2
- B. Amazon API Gateway and AWS Lambda
- C. Amazon CloudFront and Amazon S3
- D. Amazon CloudFront and AWS Elastic Beanstalk

Answer: C

NEW QUESTION 203

- (Exam Topic 3)

A developer is creating an event handling system. To handle messages asynchronously, the developer created a standard Amazon SQS queue. Quality assurance testing reveals that some events were processed multiple times.

What is the recommended way to ensure the events are not processed more than once?

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

Answer: C

NEW QUESTION 208

- (Exam Topic 3)

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

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

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

Which action will resolve this login issue?

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

Answer: A

NEW QUESTION 210

- (Exam Topic 3)

A developer implemented a static website hosted in Amazon S3 that makes web service requests hosted in Amazon API Gateway AWS Lambda. The site is showing an error that reads "No 'Access-Control-Allow-Origin' header is present on the requested resource Origin 'null' is therefore not allowed access " What should the developer do to resolve this issue?

- A. Enable cross-origin resource sharing (CORS) on the S3 bucket
- B. Enable cross-origin resource sharing (CORS) for the method in API Gateway
- C. Add the Access-Control-Request-Method header to the request
- D. Add the Access-Control-Request-Headers header to the request

Answer: B

Explanation:

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

NEW QUESTION 212

- (Exam Topic 3)

A developer is building an application using an Amazon API Gateway REST API backed by an AWS Lambda function that interacts with an Amazon DynamoDB table During testing, the developer observes high latency when making requests to the API How can the developer evaluate the end-to-end latency and identify performance bottlenecks?

- A. Enable AWS CloudTrail logging and use the logs to map each latency and bottleneck
- B. Enable and configure AWS X-Ray tracing on API Gateway and the Lambda function Use X-Ray to trace and analyze user requests
- C. Enable Amazon CloudWatch Logs for the Lambda function Enable execution logs for API Gateway to view and analyze user request logs.
- D. Enable VPC Flow Logs to capture and analyze network traffic within the VPC

Answer: B

NEW QUESTION 215

- (Exam Topic 3)

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

```
{
  "FailedRecordCount": 1,
  "Records": [
    {
      "SequenceNumber": "21269319989900637946712965403778482371",
      "ShardId": "shardId-000000000001"
    },
    {
      "ErrorCode": "ProvisionedThroughputExceededException",
      "ErrorMessage": "Rate exceeded for shard shardId-000000000001 in
        stream exampleStreamName under account 123456789."
    },
    {
      "SequenceNumber": "21269319989999637946712965403778482985",
      "ShardId": "shardId-000000000002"
    }
  ]
}
```

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

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

Answer: AC

NEW QUESTION 217

- (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 222

- (Exam Topic 3)

An organization is using Amazon API Gateway to provide a public API called "Survey" for collecting user feedback posts about its products. The survey API has "DEV" and "PROD" stages and consists of one resource "/feedback" which allows users to retrieve/create/update single feedback posts.

A version-controlled Swagger file is used to define a new API that retrieves multiple feedback posts. To add the new API resource "/listFeedbackForProduct" the developer makes changes to the Swagger file, defines an API, uploads the file to the organization's version control system, and uses the API Gateway Import API feature to apply the changes to the Survey API. After successful import, the developer runs the tests against the DEV stage and finds that resource "/listFeedbackForProduct" is not available.

What is MOST likely the reason for resource not being available?

- A. Even though the Swagger import was successful, resource creation failed afterwards.
- B. There is a propagation delay of several minutes in creating API Gateway resources after import.
- C. The developer needs to restart the API Gateway stage after import in order to apply the changes.
- D. The developer needs to create a new deployment after import in order to deploy the changes.

Answer: A

NEW QUESTION 223

- (Exam Topic 3)

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

Where should the temporary file be stored?

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

Answer: A

NEW QUESTION 228

- (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 230

- (Exam Topic 3)

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

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

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

Answer: C

NEW QUESTION 234

- (Exam Topic 3)

A developer is working on an ecommerce website. The developer wants to review server logs without logging in to each of the application servers individually. The website runs on multiple Amazon EC2 instances, is written in Python, and needs to be highly available. How can the developer update the application to meet these requirements with MINIMUM changes?

- A. Rewrite the application to be cloud native and to run on AWS Lambda where the logs can be reviewed in Amazon CloudWatch.
- B. Set up centralized logging by using Amazon Elasticsearch Service (Amazon ES), Logstash, and Kibana
- C. Scale down the application to one larger EC2 instance where only one instance is recording logs.
- D. Install the unified Amazon CloudWatch agent on the EC2 instance
- E. Configure the agent to push the application logs to CloudWatch.

Answer: D

NEW QUESTION 237

- (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 240

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 243

- (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 245

- (Exam Topic 3)

A company is launching a new web application in the AWS Cloud. The company's development team is using AWS Elastic Beanstalk for deployment and maintenance. According to the company's change management process, the development team must evaluate changes for a specific time period before completing the rollout.

Which deployment policy meets this requirement?

- A. Immutable
- B. Rolling
- C. Rolling with additional batch
- D. Traffic splitting

Answer: A

NEW QUESTION 248

- (Exam Topic 3)

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

- A. AfterInstall
- B. DownloadBundle
- C. BeforeInstall
- D. ValdateService

Answer: A

NEW QUESTION 249

- (Exam Topic 3)

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

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

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

Answer: D

NEW QUESTION 253

- (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 257

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 258

- (Exam Topic 3)

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

What is the MOST likely cause of this issue?

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

Answer: A

NEW QUESTION 262

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 265

- (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 266

- (Exam Topic 3)

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

What will be the result of this mistake?

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

Answer: A

NEW QUESTION 269

- (Exam Topic 3)

A company has a web application that uses an Amazon Cognito user pool for authentication. The company wants to create a login page with the company logo.

What should a developer do to meet these requirements?

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

Answer: A

NEW QUESTION 272

- (Exam Topic 3)

A developer is designing a distributed application built using a microservices architect spanning multiple AWS accounts. The company's operations team wants to analyze and debug application issues from a centralized account.

How can the developer meet these requirements?

- A. Use an Amazon X-Ray agent with role assumption on to publish data into the centralized account.
- B. Use Amazon X-Ray and create a new IAM user to publish the access keys into the centralized account.
- C. Use VPC Flow Logs to collect application logs across different accounts.
- D. Enable AWS CloudTrail to publish the trails in an Amazon S3 bucket in the centralized account.

Answer: A

NEW QUESTION 273

- (Exam Topic 3)

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

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

How should a developer modify the application to improve performance?

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

Answer: A

NEW QUESTION 275

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 277

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 279

- (Exam Topic 3)

A video-hosting website has two types of members: those who pay a fee, and those who do not. Each video upload places a message in Amazon SQS. A fleet of Amazon EC2 instances polls Amazon SQS and processes each video.

The developer needs to ensure that the videos uploaded by the paying members are processed first. How can the developer meet this requirement?

- A. Create two SQS queues: one for paying members, and one for non-paying members. Poll the paying member queue first and then poll the non-paying member queue.
- B. Use SQS to set priorities on individual items within a single queue: give the paying members' videos the highest priority.
- C. Use SQS to set priorities on individual items within a single queue and use Amazon SNS to encode the videos.
- D. Create two Amazon SNS topics: one for paying members and one for non-paying members. Use SNS topic subscription priorities to differentiate between the two types of members.

Answer: B

NEW QUESTION 284

- (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 287

- (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 289

- (Exam Topic 3)

A company has three AWS Lambda functions that are written in Node.js. The Lambda functions include a mix of custom code and open-source modules. When bugs are occasionally detected in the open-source modules, all three Lambda functions must be patched. What is the MOST operationally efficient solution to deploy a patched open-source library for all three Lambda functions?

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

Answer: D

NEW QUESTION 291

- (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 296

- (Exam Topic 3)

A company's e-commerce website is experiencing massive traffic spikes, which are causing performance problems in the company database. Users are reporting that accessing the website takes a long time. A developer wants to implement a caching layer using Amazon ElastiCache. The website is required to be responsive no matter which product a user views, and the updates to product information and prices must be strongly consistent.

- A. Which cache writing policy will satisfy these requirements?
- B. Write to the cache directly and sync the backend at a later time.
- C. Write to the backend first and wait for the cache to expire.
- D. Write to the cache and the backend at the same time.
- E. Write to the backend first and invalidate the cache.

Answer: E

NEW QUESTION 297

- (Exam Topic 3)

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

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

Answer: BE

NEW QUESTION 301

- (Exam Topic 3)

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

What is the MOST likely cause of this error?

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

Answer: D

NEW QUESTION 302

- (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 304

- (Exam Topic 3)

A developer at a company writes an AWS CloudFormation template. The template refers to subnets that were created by a separate AWS CloudFormation template that the company's network team wrote. When the developer attempts to launch the stack for the first time, the launch fails.

Which template coding mistakes could have caused this failure? (Select TWO.)

- A. The developer's template does not use the Ref intrinsic function to refer to the subnets
- B. The developer's template does not use the ImportValue intrinsic function to refer to the subnets
- C. The Mappings section of the developer's template does not refer to the subnets.
- D. The network team's template does not export the subnets in the Outputs section
- E. The network team's template does not export the subnets in the Mappings section

Answer: BD

NEW QUESTION 307

- (Exam Topic 3)

A developer is troubleshooting a three-tier application, which is deployed on Amazon EC2 instances. There is a connectivity problem between the application servers and the database servers.

Which AWS services or tools should be used to identify the faulty component? (Select TWO.)

- A. AWS CloudTrail.
- B. AWS Trusted Advisor
- C. Amazon VPC Flow Logs
- D. Network access control lists
- E. AWS Config rules

Answer: CD

NEW QUESTION 308

- (Exam Topic 3)

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

How should the developer reference the parameter that contains the database hostname?

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

Answer: C

NEW QUESTION 313

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 315

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 318

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 322

- (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 326

- (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 331

- (Exam Topic 3)

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

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

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

Answer: B

NEW QUESTION 334

- (Exam Topic 3)

A developer is working with a Docker application that needs to be quickly deployed using AWS without changing the infrastructure or configuring health checks. The application should be configured so that changes and updates can be made automatically without any downtime. Which solution will meet these requirements?

- A. Use AWS Elastic Beanstalk for application deployment and select an all-at-once update policy.
- B. Use AWS Elastic Beanstalk for application deployment and select a rolling deployment policy.
- C. Deploy the Docker container on an Amazon EC2 instance in an Auto Scaling group and configure a health check on the EC2 instance
- D. Deploy the Docker container using AWS Lambda and enable Amazon CloudWatch monitoring

Answer: A

NEW QUESTION 338

- (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 function
- G. View the traces from the X-Ray service.

Answer: D

NEW QUESTION 341

- (Exam Topic 3)

A three-tier application hosted on AWS uses Amazon RDS for MySQL as its database. A developer must ensure the database credentials are stored and accessed securely.

What is the MOST secure way for the developer to achieve this?

- A. Store the credentials in a configuration file and commit it to the GIT repository.
- B. Store the credentials in AWS Secrets Manager and enable automatic secret rotation.
- C. Store the credentials using Amazon RDS and enable automatic rotation
- D. Store the credentials in code and handle credentials rotation within the application.

Answer: A

NEW QUESTION 343

- (Exam Topic 3)

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 347

- (Exam Topic 3)

A developer is changing the configuration for a CPU-intensive AWS Lambda function that runs once an hour. The function usually takes 45 seconds to run, but sometimes the run time is up to 1 minute. The timeout parameter is set to 3 minutes, and all other parameters are set to default.

The developer needs to optimize the run time of this function. Which solution will meet this requirement?

- A. Redeploy the function within the default VPC
- B. Increase the function's memory.
- C. Redeploy the function by using Lambda layers
- D. Increase the function's reserved concurrency

Answer: B

NEW QUESTION 349

- (Exam Topic 3)

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

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

Answer: CE

NEW QUESTION 353

- (Exam Topic 3)

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

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

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

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

Answer: C

NEW QUESTION 356

- (Exam Topic 3)

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

How should the developer configure AWS SAM to grant the necessary read privilege to the S3 bucket?

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

Answer: D

NEW QUESTION 360

- (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 365

- (Exam Topic 3)

A developer has designed a customer-facing application that is running on an Amazon EC2 instance. The application logs every request made to it. The application usually runs seamlessly, but a spike in traffic generates several logs that cause the disk to fill up and eventually run out of memory. Company policy requires logs to be centralized for analysis.

Which long-term solution should the developer employ to prevent the issue from reoccurring?

- A. Install the Amazon CloudWatch agent on the instance to send the logs to CloudWatch
- B. Delete the logs from the instance once they are sent to CloudWatch.
- C. Enable AWS Auto Scaling on Amazon Elastic Block Store (Amazon EBS) to automatically add volumes to the instance when it reaches a specified threshold.
- D. Enable AWS Auto Scaling on Amazon Elastic Block Store (Amazon EBS) to automatically add volume to the instance when it reaches a specified threshold.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to pull the logs from the instance. Configure the rule to delete the logs they have been pulled.

Answer: D

NEW QUESTION 367

- (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 371

- (Exam Topic 3)

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

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

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

Answer: B

NEW QUESTION 376

- (Exam Topic 3)

When using the AWS Encryption SDK how does the developer keep track of the data encryption keys used to encrypt data?

- A. The developer must manually Keep track of the data encryption keys used for each data object
- B. The SDK encrypts the data encryption key and stores it (encrypted) as part of the returned ciphertext
- C. The SDK stores the data encryption keys automatically in Amazon S3
- D. The data encryption key is stored in the userdata for the EC2 instance

Answer: B

NEW QUESTION 377

- (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 380

- (Exam Topic 2)

A company has a two-tier application running on an Amazon EC2 server that handles all of its AWS based

e-commerce activity During peak times, the backend servers that process orders are overloaded with requests. This results in some orders failing to process. A

developer needs to create a solution that will re-factor the application.

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

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

Answer: CD

NEW QUESTION 384

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

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

100% Actual & Verified — Instant Download, Please Click
[Order The DVA-C01 Practice Test Here](#)