

DOP-C02 Exam Format, DOP-C02 Dumps Free



BTW, DOWNLOAD part of Exam4Tests DOP-C02 dumps from Cloud Storage: <https://drive.google.com/open?id=1V79peMb1oPaU5s15DlqwIO28QZp0Adls>

Exam4Tests is a wonderful study platform that can transform your effective diligence into your best rewards. By years of diligent work, our experts have collected the frequent-tested knowledge into our DOP-C02 exam materials for your reference. So our DOP-C02 Practice Questions are triumph of their endeavor. I can say that no one can know the DOP-C02 study guide better than them and our quality of the DOP-C02 learning quiz is the best.

Most of the experts in our company have been studying in the professional field for many years and have accumulated much experience in our DOP-C02 practice questions. Our company is considerably cautious in the selection of talent and always hires employees with store of specialized knowledge and skills. All the members of our experts and working staff maintain a high sense of responsibility, which is why there are so many people choose our DOP-C02 Exam Materials and to be our long-term partner.

>> DOP-C02 Exam Format <<

Quiz DOP-C02 - Reliable AWS Certified DevOps Engineer - Professional Exam Format

After the user has purchased our DOP-C02 learning materials, we will discover in the course of use that our product design is extremely scientific and reasonable. Details determine success or failure, so our every detail is strictly controlled. For example, our learning material's Windows Software page is clearly, our DOP-C02 Learning material interface is simple and beautiful. There are no additional ads to disturb the user to use the DOP-C02 qualification question. Once you have submitted your practice time, DOP-C02 study tool system will automatically complete your operation.

Amazon DOP-C02 (AWS Certified DevOps Engineer - Professional) certification exam is designed for professionals who have experience in the field of DevOps and are looking to validate their skills and knowledge. AWS Certified DevOps Engineer - Professional certification is targeted towards individuals who have a deep understanding of the AWS cloud platform and its services.

To prepare for the Amazon DOP-C02 certification exam, candidates can take advantage of various resources provided by AWS, such as training courses, practice exams, and sample questions. Candidates can also leverage their experience with AWS services and DevOps methodologies to prepare for the exam. It is recommended that candidates have at least two years of experience with

AWS services and one year of experience with DevOps practices before attempting the certification exam

Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q291-Q296):

NEW QUESTION # 291

A company has configured an Amazon S3 event source on an AWS Lambda function. The company needs the Lambda function to run when a new object is created or an existing object is modified. In a particular S3 bucket, the Lambda function will use the S3 bucket name and the S3 object key of the incoming event to read the contents of the created or modified S3 object. The Lambda function will parse the contents and save the parsed contents to an Amazon DynamoDB table.

The Lambda function's execution role has permissions to read from the S3 bucket and to write to the DynamoDB table. During testing, a DevOps engineer discovers that the Lambda function does not run when objects are added to the S3 bucket or when existing objects are modified.

Which solution will resolve this problem?

- A. Increase the memory of the Lambda function to give the function the ability to process large files from the S3 bucket.
- B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an OnFailure destination for the Lambda function
- **C. Create a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket**
- D. Provision space in the /tmp folder of the Lambda function to give the function the ability to process large files from the S3 bucket

Answer: C

Explanation:

Explanation

Option A is incorrect because increasing the memory of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source.

Increasing the memory of the Lambda function might improve its performance or reduce its execution time, but it does not affect its invocation. Moreover, increasing the memory of the Lambda function might incur higher costs, as Lambda charges based on the amount of memory allocated to the function.

Option B is correct because creating a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket is a necessary step to configure an S3 event source. A resource policy is a JSON document that defines who can access a Lambda resource and under what conditions. By granting Amazon S3 permission to invoke the Lambda function, the company ensures that the Lambda function runs when a new object is created or an existing object is modified in the S3 bucket.

Option C is incorrect because configuring an Amazon Simple Queue Service (Amazon SQS) queue as an On-Failure destination for the Lambda function does not help with triggering the Lambda function.

An On-Failure destination is a feature that allows Lambda to send events to another service, such as SQS or Amazon Simple Notification Service (Amazon SNS), when a function invocation fails.

However, this feature only applies to asynchronous invocations, and S3 event sources use synchronous invocations. Therefore, configuring an SQS queue as an On-Failure destination would have no effect on the problem.

Option D is incorrect because provisioning space in the /tmp folder of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source. Provisioning space in the /tmp folder of the Lambda function might help with processing large files from the S3 bucket, as it provides temporary storage for up to 512 MB of data. However, it does not affect the invocation of the Lambda function.

References:

[Using AWS Lambda with Amazon S3](#)

[Lambda resource access permissions](#)

[AWS Lambda destinations](#)

[\[AWS Lambda file system\]](#)

NEW QUESTION # 292

A software team is using AWS CodePipeline to automate its Java application release pipeline. The pipeline consists of a source stage, then a build stage, and then a deploy stage. Each stage contains a single action that has a runOrder value of 1.

The team wants to integrate unit tests into the existing release pipeline. The team needs a solution that deploys only the code changes that pass all unit tests.

Which solution will meet these requirements?

- A. Modify the build stage. Add a test action that has a runOrder value of 1. Use AWS CodeDeploy as the action provider to run unit tests.
- B. **Modify the build stage Add a test action that has a runOrder value of 2 Use AWS CodeBuild as the action provider to run unit tests**
- C. Modify the deploy stage Add a test action that has a runOrder value of 1 Use AWS CodeDeploy as the action provider to run unit tests
- D. Modify the deploy stage Add a test action that has a runOrder value of 2 Use AWS CodeBuild as the action provider to run unit tests

Answer: B

Explanation:

* Modify the Build Stage to Add a Test Action with a RunOrder Value of 2:

The build stage in AWS CodePipeline can have multiple actions. By adding a test action with a runOrder value of 2, the test action will execute after the initial build action completes.

* Use AWS CodeBuild as the Action Provider to Run Unit Tests:

AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages.

Using CodeBuild to run unit tests ensures that the tests are executed in a controlled environment and that only the code changes that pass the unit tests proceed to the deploy stage.

Example configuration in CodePipeline:

```
{
  "name": "BuildStage",
  "actions": [
    {
      "name": "Build",
      "actionTypeId": {
        "category": "Build",
        "owner": "AWS",
        "provider": "CodeBuild",
        "version": "1"
      },
      "runOrder": 1
    },
    {
      "name": "Test",
      "actionTypeId": {
        "category": "Test",
        "owner": "AWS",
        "provider": "CodeBuild",
        "version": "1"
      },
      "runOrder": 2
    }
  ]
}
```

By integrating the unit tests into the build stage and ensuring they run after the build process, the pipeline guarantees that only code changes passing all unit tests are deployed.

Reference:

[AWS CodePipeline](#)

[AWS CodeBuild](#)

[Using CodeBuild with CodePipeline](#)

NEW QUESTION # 293

A production account has a requirement that any Amazon EC2 instance that has been logged in to manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with the Amazon CloudWatch Logs agent configured.

How can this process be automated?

- A. Create an Amazon CloudWatch alarm that will be invoked by the login event. Send the notification to an Amazon Simple Notification Service (Amazon SNS) topic that the operations team is subscribed to, and have them terminate the EC2

instance within 24 hours.

- B. Create a CloudWatch Logs subscription to an AWS Step Functions application. Configure an AWS Lambda function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned. Create an Amazon EventBridge rule to invoke a second Lambda function once a day that will terminate all instances with this tag.
- C. Create a CloudWatch Logs subscription to an AWS Lambda function. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned. Create an Amazon EventBridge rule to invoke a daily Lambda function that terminates all instances with this tag.
- D. Create an Amazon CloudWatch alarm that will be invoked by the login event. Configure the alarm to send to an Amazon Simple Queue Service (Amazon SQS) queue. Use a group of worker instances to process messages from the queue, which then schedules an Amazon EventBridge rule to be invoked.

Answer: C

NEW QUESTION # 294

A company's application uses a fleet of Amazon EC2 On-Demand Instances to analyze and process data. The EC2 instances are in an Auto Scaling group. The Auto Scaling group is a target group for an Application Load Balancer (ALB). The application analyzes critical data that cannot tolerate interruption. The application also analyzes noncritical data that can withstand interruption.

The critical data analysis requires quick scalability in response to real-time application demand. The noncritical data analysis involves memory consumption. A DevOps engineer must implement a solution that reduces scale-out latency for the critical data. The solution also must process the noncritical data.

Which combination of steps will meet these requirements? (Select TWO.)

- A. For the noncritical data, create a second Auto Scaling group that uses a launch template. Configure the launch template to install the unified Amazon CloudWatch agent and to configure the CloudWatch agent with a custom memory utilization metric. Use Spot Instances. Add the new Auto Scaling group as the target group for the ALB. Modify the application to use two target groups for critical data and noncritical data.
- B. For the critical data, modify the existing Auto Scaling group. Create a warm pool instance in the stopped state. Define the warm pool size. Create a new version of the launch template that has detailed monitoring enabled. Use On-Demand Instances.
- C. For the noncritical data, create a second Auto Scaling group. Choose the predefined memory utilization metric type for the target tracking scaling policy. Use Spot Instances. Add the new Auto Scaling group as the target group for the ALB. Modify the application to use two target groups for critical data and noncritical data.
- D. For the critical data, modify the existing Auto Scaling group. Create a lifecycle hook to ensure that bootstrap scripts are completed successfully. Ensure that the application on the instances is ready to accept traffic before the instances are registered. Create a new version of the launch template that has detailed monitoring enabled.
- E. For the critical data, modify the existing Auto Scaling group. Create a warm pool instance in the stopped state. Define the warm pool size. Create a new version of the launch template that has detailed monitoring enabled. use Spot Instances.

Answer: A,B

Explanation:

Explanation

For the critical data, using a warm pool¹ can reduce the scale-out latency by having pre-initialized EC2 instances ready to serve the application traffic. Using On-Demand Instances can ensure that the instances are always available and not interrupted by Spot interruptions².

For the noncritical data, using a second Auto Scaling group with Spot Instances can reduce the cost and leverage the unused capacity of EC2³. Using a launch template with the CloudWatch agent⁴ can enable the collection of memory utilization metrics, which can be used to scale the group based on the memory demand. Adding the second group as a target group for the ALB and modifying the application to use two target groups can enable routing the traffic based on the data type.

References: 1: Warm pools for Amazon EC2 Auto Scaling 2: Amazon EC2 On-Demand Capacity Reservations 3: Amazon EC2 Spot Instances 4: Metrics collected by the CloudWatch agent

NEW QUESTION # 295

A company has a new AWS account that teams will use to deploy various applications. The teams will create many Amazon S3 buckets for application- specific purposes and to store AWS CloudTrail logs. The company has enabled Amazon Macie for the account.

A DevOps engineer needs to optimize the Macie costs for the account without compromising the account's functionality.

Which solutions will meet these requirements? (Select TWO.)

- A. Configure discovery jobs to include S3 objects that are tagged as production only.
- B. Configure scheduled daily discovery jobs for all S3 buckets in the account.
- C. **Configure discovery jobs to include S3 objects based on the last modified criterion.**
- D. Exclude S3 buckets that have public read access from automated discovery.
- E. Exclude S3 buckets that contain CloudTrail logs from automated discovery.

Answer: C,E

Explanation:

Explanation
To optimize the Macie costs for the account without compromising the account's functionality, the DevOps engineer needs to exclude S3 buckets that do not contain sensitive data from automated discovery. S3 buckets that contain CloudTrail logs are unlikely to have sensitive data, and Macie charges for scanning and monitoring data in S3 buckets. Therefore, excluding S3 buckets that contain CloudTrail logs from automated discovery can reduce Macie costs. Similarly, configuring discovery jobs to include S3 objects based on the last modified criterion can also reduce Macie costs, as it will only scan and monitor new or updated objects, rather than all objects in the bucket.

NEW QUESTION # 296

• • • • •

Our experts are constantly looking for creative ways to immortalize our DOP-C02 actual exam in this line. Their masterpieces are instrumental to offer help and improve your performance in the real exam. Being dedicated to these practice materials painstakingly and pooling useful points into our DOP-C02 Exam Materials with perfect arrangement and scientific compilation of messages, our DOP-C02 practice materials can propel the exam candidates to practice with efficiency.

DOP-C02 Dumps Free: <https://www.exam4tests.com/DOP-C02-valid-braindumps.html>

DOWNLOAD the newest Exam4Tests DOP-C02 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1V79peMb1oPaU5s15DlwI028QZp0Adls>

