AWS Certified DevOps Engineer - Professional exam vce torrent & DOP-C02 pdf dumps & AWS Certified DevOps Engineer - Professional valid study prep

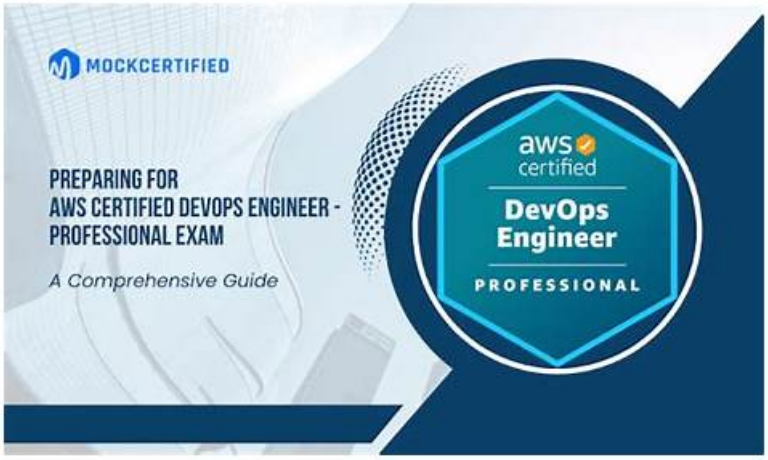

2025 Latest ActualtestPDF DOP-C02 PDF Dumps and DOP-C02 Exam Engine Free Share: https://drive.google.com/open?id=1DJ9w_toEsG_o0IwfUmdiliEpPo0xElve

ActualtestPDF's study material is available in three different formats. The reason we have introduced three formats of the AWS Certified DevOps Engineer - Professional (DOP-C02) practice material is to meet the learning needs of every student. Some candidates prefer DOP-C02 practice exams and some want Real DOP-C02 Questions due to a shortage of time. At ActualtestPDF, we meet the needs of both types of aspirants. We have Amazon DOP-C02 PDF format, a web-based practice exam, and AWS Certified DevOps Engineer - Professional (DOP-C02) desktop practice test software.

With the development of science and technology the internet in our daily life is playing a more and more important role! IT workers become high-salary people. Amazon certifications become hot vocational qualification certificate. ActualtestPDF offers the best DOP-C02 Guide Torrent files to help people clear exams and realize their idea better. We are engaged in this field more than 8 years. If you have dream in this field, our valid DOP-C02 guide torrent files will be a good chance for you.

>> DOP-C02 Interactive Course <<

Valid Braindumps DOP-C02 Ebook, VCE DOP-C02 Dumps

Three versions for DOP-C02 exam cram are available, and you can choose the most suitable one according to your own needs. DOP-C02 Online test engine supports all web browsers, and you can also have offline practice. One of the most outstanding features of DOP-C02 Online test engine is that it has testing history and performance review, and you can have a general review of what you have learnt through this version. DOP-C02 Soft test engine supports MS operating system as well as stimulates real exam environment, therefore it can build up your confidence. DOP-C02 PDF version is printable, and you can study anytime.

Amazon DOP-C02 (AWS Certified DevOps Engineer - Professional) certification exam is a highly sought after certification that validates the skills and knowledge required to manage and deploy applications on the AWS platform. AWS Certified DevOps Engineer - Professional certification is designed for DevOps engineers who have experience in developing, provisioning, operating and managing applications on the AWS platform. DOP-C02 Exam Tests the candidate's ability to design, deploy, manage, and maintain AWS-based applications using DevOps practices and principles.

Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q362-Q367):

NEW QUESTION # 362
AnyCompany is using AWS Organizations to create and manage multiple AWS accounts AnyCompany recently acquired a smaller company, Example Corp. During the acquisition process, Example Corp's single AWS account joined AnyCompany's management account through an Organizations invitation. AnyCompany moved the new member account under an OU that is dedicated to

Example Corp.
AnyCompany's DevOps eng*neer has an IAM user that assumes a role that is named OrganizationAccountAccessRole to access member accounts. This role is configured with a full access policy When the DevOps engineer tries to use the AWS Management Console to assume the role in Example Corp's new member account, the DevOps engineer receives the following error message "Invalid information in one or more fields. Check your information or contact your administrator." Which solution will give the DevOps engineer access to the new member account?

- A. In the new member account, create a new IAM role that is named OrganizationAccountAccessRole. Attach the AdmInistratorAccess AVVS managed policy to the role. In the role's trust policy, grant the management account permission to assume the role.
- B. In the management account, create a new SCR In the SCP, grant the DevOps engineer's IAM user full access to all resources in the new member account. Attach the SCP to the OU that contains the new member account,
- C. In the management account, grant the DevOps engineer's IAM user permission to assume the OrganzatlonAccountAccessR01e IAM role in the new member account.
- D. In the new member account edit the trust policy for the Organ zationAccountAccessRole IAM role. Grant the management account permission to assume the role.

Answer: A

Explanation:
The problem is that the DevOps engineer cannot assume the OrganizationAccountAccessRole IAM role in the new member account that joined AnyCompany's management account through an Organizations invitation. The solution is to create a new IAM role with the same name and trust policy in the new member account.
Option A is incorrect, as it does not address the root cause of the error. The DevOps engineer's IAM user already has permission to assume the OrganizationAccountAccessRole IAM role in any member account, as this is the default role name that AWS Organizations creates when a new account joins an organization. The error occurs because the new member account does not have this role, as it was not created by AWS Organizations.
Option B is incorrect, as it does not address the root cause of the error. An SCP is a policy that defines the maximum permissions for account members of an organization or organizational unit (OU). An SCP does not grant permissions to IAM users or roles, but rather limits the permissions that identity-based policies or resource-based policies grant to them. An SCP also does not affect how IAM roles are assumed by other principals.
Option C is correct, as it addresses the root cause of the error. By creating a new IAM role with the same name and trust policy as the OrganizationAccountAccessRole IAM role in the new member account, the DevOps engineer can assume this role and access the account. The new role should have the AdministratorAccess AWS managed policy attached, which grants full access to all AWS resources in the account. The trust policy should allow the management account to assume the role, which can be done by specifying the management account ID as a principal in the policy statement.
Option D is incorrect, as it assumes that the new member account already has the OrganizationAccountAccessRole IAM role, which is not true. The new member account does not have this role, as it was not created by AWS Organizations. Editing the trust policy of a non-existent role will not solve the problem.

NEW QUESTION # 363
A company has configured Amazon RDS storage autoscaling for its RDS DB instances. A DevOps team needs to visualize the autoscaling events on an Amazon CloudWatch dashboard. Which solution will meet this requirement?

- A. Create an Amazon EventBridge rule that reacts to RDS storage autoscaling events (rom the RDS events. Create a CloudWatch alarm. Configure the EventBridge rule to change the status of the CloudWatch alarm. Visualize the alarm status by using the CloudWatch dashboard.
- B. Create a trail by using AWS CloudTrail with data events configured. Configure the trail to send the data events to Amazon CloudWatch Logs. Create a metric filter in CloudWatch Logs to match the RDS storage autoscaling events. Visualize the metric filter by using the CloudWatch dashboard.
- C. Create an Amazon EventBridge rule that reacts to RDS storage autoscaling events from RDS events.
  Create an AWS Lambda function that publishes a CloudWatch custom metric. Configure the EventBridge rule to invoke the Lambda function. Visualize the custom metric by using the CloudWatch dashboard.
- D. Create a trail by using AWS CloudTrail with management events configured. Configure the trail to send the management events to Amazon CloudWatch Logs. Create a metric filter in CloudWatch Logs to match the RDS storage autoscaling events. Visualize the metric filter by using the CloudWatch dashboard.

Answer: C

Explanation:

Step 1: Reacting to RDS Storage Autoscaling Events Using Amazon EventBridgeAmazon RDS emits events when storage autoscaling occurs. To visualize these events in a CloudWatch dashboard, you can create an EventBridge rule that listens for these specific autoscaling events.
Action: Create an EventBridge rule that reacts to RDS storage autoscaling events from the RDS event stream.
Why: EventBridge allows you to listen to RDS events and route them to specific AWS services for processing.
Step 2: Creating a Custom CloudWatch Metric via LambdaOnce the EventBridge rule detects a storage autoscaling event, you can use a Lambda function to publish a custom metric to CloudWatch. This metric can then be visualized in a CloudWatch dashboard.
Action: Use a Lambda function to publish custom metrics to CloudWatch based on the RDS storage autoscaling events.
Why: Custom metrics allow you to track specific events like autoscaling and visualize them easily on a CloudWatch dashboard.
Reference: AWS documentation on Publishing Custom Metrics to CloudWatch.
This corresponds to Option A: Create an Amazon EventBridge rule that reacts to RDS storage autoscaling events from RDS events. Create an AWS Lambda function that publishes a CloudWatch custom metric.
Configure the EventBridge rule to invoke the Lambda function. Visualize the custom metric by using the CloudWatch dashboard.

NEW QUESTION # 364
A company uses an AWS CodeCommit repository to store its source code and corresponding unit tests. The company has configured an AWS CodePipeline pipeline that includes an AWS CodeBuild project that runs when code is merged to the main branch of the repository.
The company wants the CodeBuild project to run the unit tests. If the unit tests pass, the CodeBuild project must tag the most recent commit.
How should the company configure the CodeBuild project to meet these requirements?

- A. Configure the CodeBuild project to use native Git to clone the CodeCommit repository. Configure the project to run the unit tests. Configure the project to use native Git to create a tag and to push the Git tag to the repository if the code passes the unit tests.
- B. Configure the CodeBuild project to use native Git to clone the CodeCommit repository. Configure the project to run the unit tests. Configure the project to use AWS CLI commands to create a new repository tag in the repository if the code passes the unit tests.
- C. Configure the CodeBuild project to use AWS CLI commands to copy the code from the CodeCommit repository. Configure the project lo run the unit tests. Configure the project to use AWS CLI commands to create a new Git tag in the repository if the code passes the unit tests.
- D. Configure the CodeBuild project to use AWS CLI commands to copy the code from the CodeCommit repository. Configure the project to run the unit tests. Configure the project to use AWS CLI commands to create a new repository tag in the repository if the code passes the unit tests.

Answer: A

Explanation:
Step 1: Using Native Git in CodeBuild
To meet the requirement of running unit tests and tagging the most recent commit if the tests pass, the CodeBuild project should be configured to use native Git to clone the CodeCommit repository. Native Git support allows full functionality for managing the repository, including the ability to create and push tags.
Action: Configure the CodeBuild project to use native Git to clone the repository and run the tests.
Why: Using native Git provides flexibility for managing tags and other repository operations after the tests are successfully executed.
Step 2: Tagging the Most Recent Commit
Once the unit tests pass, the CodeBuild project can use native Git to create a tag for the most recent commit and push that tag to the repository. This ensures that the tagged commit is linked to the test results.
Action: Configure the project to use native Git to create and push a tag to the repository if the tests pass.
Why: This ensures the correct commit is tagged automatically, streamlining the workflow.
Reference:
This corresponds to Option A: Configure the CodeBuild project to use native Git to clone the CodeCommit repository. Configure the project to run the unit tests. Configure the project to use native Git to create a tag and to push the Git tag to the repository if the code passes the unit tests.

NEW QUESTION # 365
A company manages an application that stores logs in Amazon CloudWatch Logs. The company wants to archive the logs to an Amazon S3 bucket. Logs are rarely accessed after 90 days and must be retained for 10 years.
Which combination of steps should a DevOps engineer take to meet these requirements? (Select TWO.)

- A. Configure the S3 bucket lifecycle policy to transition logs to S3 Glacier Instant Retrieval after 90 days and to expire logs after 3,650 days.
- B. Configure a CloudWatch Logs subscription filter to use AWS Glue to transfer all logs to an S3 bucket.
- C. Configure the S3 bucket lifecycle policy to transition logs to Reduced Redundancy after 90 days and to expire logs after 3,650 days.
- D. Configure a CloudWatch Logs subscription filter to stream all logs to an S3 bucket.
- E. Configure a CloudWatch Logs subscription filter to use Amazon Data Firehose to stream all logs to an S3 bucket.

Answer: A,E

Explanation:
To archive CloudWatch Logs to S3 with long-term retention, you need:
* a streaming mechanism to move data from CloudWatch Logs to S3, and
* an S3 Lifecycle policy to handle tiering and expiration.
Option B uses a CloudWatch Logs subscription filter that targets Amazon Data Firehose. Firehose provides managed, scalable streaming from CloudWatch Logs to S3 with optional buffering and transformation. This is the AWS-recommended pattern for exporting continuous log streams with minimal operational overhead.
Option C is not valid because CloudWatch Logs subscription filters cannot directly target S3; they must target Kinesis, Firehose, or Lambda.
Once logs are in S3, the company wants to keep them rarely accessed after 90 days and retained for 10 years.
Option D configures an S3 lifecycle policy to transition logs to S3 Glacier Instant Retrieval after 90 days, which is a low-cost archival tier with relatively fast access, and to expire (delete) logs after 3,650 days (10 years). This precisely matches the retention requirement.
Option E uses Reduced Redundancy, which is a legacy storage class and not optimized for long-term archival.
Therefore, the correct combination is B and D.

NEW QUESTION # 366
A company recently deployed its web application on AWS. The company is preparing for a large-scale sales event and must ensure that the web application can scale to meet the demand The application's frontend infrastructure includes an Amazon CloudFront distribution that has an Amazon S3 bucket as an origin. The backend infrastructure includes an Amazon API Gateway API. several AWS Lambda functions, and an Amazon Aurora DB cluster The company's DevOps engineer conducts a load test and identifies that the Lambda functions can fulfill the peak number of requests However, the DevOps engineer notices request latency during the initial burst of requests Most of the requests to the Lambda functions produce queries to the database A large portion of the invocation time is used to establish database connections Which combination of steps will provide the application with the required scalability? (Select TWO)

- A. Configure a higher reserved concurrency for the Lambda functions.
- B. Use Amazon RDS Proxy to create a proxy for the Aurora database Update the Lambda functions to use the proxy endpoints for database connections.
- C. Configure a higher provisioned concurrency for the Lambda functions
- D. Convert the DB cluster to an Aurora global database Add additional Aurora Replicas in AWS Regions based on the locations of the company's customers.
- E. Refactor the Lambda Functions Move the code blocks that initialize database connections into the function handlers.

Answer: B,C

Explanation:
Explanation
The correct answer is B and E. Configuring a higher provisioned concurrency for the Lambda functions will ensure that the functions are ready to respond to the initial burst of requests without any cold start latency.
Using Amazon RDS Proxy to create a proxy for the Aurora database will enable the Lambda functions to reuse existing database connections and reduce the overhead of establishing new ones. This will also improve the scalability and availability of the database by managing the connection pool size and handling failovers.
Option A is incorrect because reserved concurrency only limits the number of concurrent executions for a function, not pre-warms them. Option C is incorrect because converting the DB cluster to an Aurora global database will not address the issue of database connection latency, and may introduce additional costs and complexity. Option D is incorrect because moving the code blocks that initialize database connections into the function handlers will not improve the performance or scalability of the Lambda functions, and may actually worsen the cold start latency. References:
* AWS Lambda Provisioned Concurrency
* Using Amazon RDS Proxy with AWS Lambda

* Certified DevOps Engineer - Professional (DOP-C02) Study Guide (page 173)

NEW QUESTION # 367
......

The ActualtestPDF offers valid, updated, and real AWS Certified DevOps Engineer - Professional DOP-C02 exam practice questions that perfectly and quickly prepare the DOP-C02 exam candidates. You can easily pass the challenging AWS Certified DevOps Engineer - Professional DOP-C02 Certification Exam. DOP-C02 exam practice test questions you will get everything that you need to learn, prepare and pass the valuable DOP-C02 certification with good scores.

Valid Braindumps DOP-C02 Ebook: https://www.actualtestpdf.com/Amazon/DOP-C02-practice-exam-dumps.html

2025 Latest ActualtestPDF DOP-C02 PDF Dumps and DOP-C02 Exam Engine Free Share: https://drive.google.com/open?id=1DJ9w_toEsG_o0IwfUmdiliEpPo0xElve