



DOP-C02^{Q&As}

AWS Certified DevOps Engineer - Professional

Pass Amazon DOP-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/dop-c02.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

A development team manually builds an artifact locally and then places it in an Amazon S3 bucket. The application has a local cache that must be cleared when a deployment occurs. The team runs a command to do this downloads the artifact from Amazon S3 and unzips the artifact to complete the deployment.

A DevOps team wants to migrate to a CI/CD process and build in checks to stop and roll back the deployment when a failure occurs. This requires the team to track the progression of the deployment.

Which combination of actions will accomplish this? (Select THREE)

- A. Allow developers to check the code into a code repository Using Amazon EventBridge on every pull into the main branch invoke an AWS Lambda function to build the artifact and store it in Amazon S3.
- B. Create a custom script to clear the cache Specify the script in the BeforeInstall lifecycle hook in the AppSpec file.
- C. Create user data for each Amazon EC2 instance that contains the clear cache script Once deployed test the application If it is not successful deploy it again.
- D. Set up AWS CodePipeline to deploy the application Allow developers to check the code into a code repository as a source for the pipeline.
- E. Use AWS CodeBuild to build the artifact and place it in Amazon S3 Use AWS CodeDeploy to deploy the artifact to Amazon EC2 instances.
- F. Use AWS Systems Manager to fetch the artifact from Amazon S3 and deploy it to all the instances.

Correct Answer: BDE

QUESTION 2

A DevOps engineer is building an application that uses an AWS Lambda function to query an Amazon Aurora MySQL DB cluster. The Lambda function performs only read queries. Amazon EventBridge events invoke the Lambda function.

As more events invoke the Lambda function each second, the database's latency increases and the database's throughput decreases. The DevOps engineer needs to improve the performance of the application.

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

- A. Use Amazon RDS Proxy to create a proxy. Connect the proxy to the Aurora cluster reader endpoint. Set a maximum connections percentage on the proxy.
- B. Implement database connection pooling inside the Lambda code. Set a maximum number of connections on the database connection pool.
- C. Implement the database connection opening outside the Lambda event handler code.
- D. Implement the database connection opening and closing inside the Lambda event handler code.
- E. Connect to the proxy endpoint from the Lambda function.
- F. Connect to the Aurora cluster endpoint from the Lambda function.



Correct Answer: ACE

To improve the performance of the application, the DevOps engineer should use Amazon RDS Proxy, implement the database connection opening outside the Lambda event handler code, and connect to the proxy endpoint from the Lambda

function.

References:

Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure¹. By using Amazon RDS

Proxy, the DevOps engineer can reduce the overhead of opening and closing connections to the database, which can improve latency and throughput². The DevOps engineer should connect the proxy to the Aurora cluster reader endpoint,

which allows read-only connections to one of the Aurora Replicas in the DB cluster³. This can help balance the load across multiple read replicas and improve performance for read-intensive workloads⁴. The DevOps engineer should

implement the database connection opening outside the Lambda event handler code, which means using a global variable to store the database connection object⁵. This can enable connection reuse across multiple invocations of the

Lambda function, which can reduce latency and improve performance.

The DevOps engineer should connect to the proxy endpoint from the Lambda function, which is a unique URL that represents the proxy. This can allow the Lambda function to access the database through the proxy, which can provide benefits such as connection pooling, load balancing, failover handling, and enhanced security.

QUESTION 3

A DevOps engineer is developing an application for a company. The application needs to persist files to Amazon S3. The application needs to upload files with different security classifications that the company defines. These classifications include confidential, private, and public. Files that have a confidential classification must not be viewable by anyone other than the user who uploaded them. The application uses the IAM role of the user to call the S3 API operations.

The DevOps engineer has modified the application to add a DataClassification tag with the value of confidential and an Owner tag with the uploading user's ID to each confidential object that is uploaded to Amazon S3.

Which set of additional steps must the DevOps engineer take to meet the company's requirements?

A. Modify the S3 bucket's ACL to grant bucket-owner-read access to the uploading user's IAM role. Create an IAM policy that grants s3:GetObject operations on the S3 bucket when aws:ResourceTag/DataClassification equals confidential, and s3:ExistingObjectTag/Owner equals \${aws:userid}. Attach the policy to the IAM roles for users who require access to the S3 bucket.

B. Modify the S3 bucket policy to allow the s3:GetObject action when aws:ResourceTag/DataClassification equals confidential, and s3:ExistingObjectTag/Owner equals \${aws:userid}. Create an IAM policy that grants s3:GetObject operations on the S3 bucket. Attach the policy to the IAM roles for users who require access to the S3 bucket.

C. Modify the S3 bucket policy to allow the s3:GetObject action when aws:ResourceTag/DataClassification equals confidential, and aws:RequestTag/Owner equals \${aws:userid}. Create an IAM policy that grants s3:GetObject operations on the S3 bucket. Attach the policy to the IAM roles for users who require access to the S3 bucket.



D. Modify the S3 bucket's ACL to grant authenticated-read access when `aws:ResourceTag/DataClassification` equals `confidential`, and `s3:ExistingObjectTag/Owner` equals `${aws:userid}`. Create an IAM policy that grants `s3:GetObject` operations on the S3 bucket. Attach the policy to the IAM roles for users who require access to the S3 bucket.

Correct Answer: B

QUESTION 4

What storage driver does Docker generally recommend that you use if it is available?

- A. zfs
- B. btrfs
- C. aufs
- D. overlay

Correct Answer: C

After you have read the storage driver overview, the next step is to choose the best storage driver for your workloads. In making this decision, there are three high-level factors to consider: If multiple storage drivers are supported in your kernel, Docker has a prioritized list of which storage driver to use if no storage driver is explicitly configured, assuming that the prerequisites for that storage driver are met: If aufs is available, default to it, because it is the oldest storage driver. However, it is not universally available.

Reference: <https://docs.docker.com/engine/userguide/storagedriver/selectadriver/>

QUESTION 5

You have just recently deployed an application on EC2 instances behind an ELB. After a couple of weeks, customers are complaining on receiving errors from the application. You want to diagnose the errors and are trying to get errors from the ELB access logs. But the ELB access logs are empty. What is the reason for this.

- A. You do not have the appropriate permissions to access the logs
- B. You do not have your CloudWatch metrics correctly configured
- C. ELB Access logs are only available for a maximum of one week
- D. Access logging is an optional feature of Elastic Load Balancing that is disabled by default

Correct Answer: D

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues. Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.



VCE & PDF

PassApply.com

<https://www.passapply.com/dop-c02.html>

2024 Latest passapply DOP-C02 PDF and VCE dumps Download

[DOP-C02 PDF Dumps](#)

[DOP-C02 Practice Test](#)

[DOP-C02 Braindumps](#)