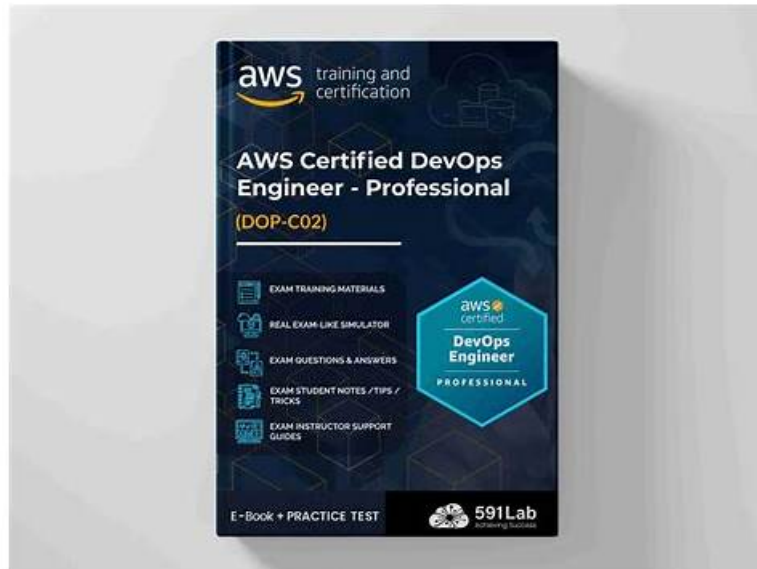


DOP-C02 Test Braindumps: AWS Certified DevOps Engineer - Professional & DOP-C02 Quiz Materials & DOP-C02 Exam Torrent



2026 Latest DumpExam DOP-C02 PDF Dumps and DOP-C02 Exam Engine Free Share: https://drive.google.com/open?id=1VITYzXONb8v7wTAuiiFOp4QNKXdfQ7_b

DOP-C02 practice test can be your optimum selection and useful tool to deal with the urgent challenge. With over a decade's striving, our DOP-C02 training materials have become the most widely-lauded and much-anticipated products in industry. We have three versions of DOP-C02 Exam Questions by modernizing innovation mechanisms and fostering a strong pool of professionals. Therefore, rest assured of full technical support from our professional elites in planning and designing DOP-C02 practice test.

We believe that you can buy our DOP-C02 demo PDF torrent without any misgivings, Firstly, we have a strong experts team who are devoted themselves to research of the technology, which ensure the high-quality of our DOP-C02 Dump guide, DumpExam offers AWS Certified DevOps Engineer - Professional DOP-C02 free Updates. It is no exaggeration to say that the value of the certification training materials is equivalent to all exam related reference books.

>> Free DOP-C02 Exam Questions <<

2026 Amazon DOP-C02 Perfect Free Exam Questions

Competition appear everywhere in modern society. There are many way to improve ourselves and learning methods of DOP-C02 exams come in different forms. Economy rejuvenation and social development carry out the blossom of technology; some DOP-C02 Learning Materials are announced which have a good quality. Certification qualification exam materials are a big industry and many companies are set up for furnish a variety of services for it.

Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q117-Q122):

NEW QUESTION # 117

A company uses AWS CodeArtifact to centrally store Python packages. The CodeArtifact repository is configured with the following repository policy.

□ A development team is building a new project in an account that is in an organization in AWS Organizations.

The development team wants to use a Python library that has already been stored in the CodeArtifact repository in the organization. The development team uses AWS CodePipeline and AWS CodeBuild to build the new application. The CodeBuild job that the development team uses to build the application is configured to run in a VPC Because of compliance requirements the VPC has no internet connectivity.

The development team creates the VPC endpoints for CodeArtifact and updates the CodeBuild buildspec yaml file. However, the development team cannot download the Python library from the repository.

Which combination of steps should a DevOps engineer take so that the development team can use Code Artifact? (Select TWO.)

- A. Create an Amazon S3 gateway endpoint Update the route tables for the subnets that are running the CodeBuild job.
- B. Specify the account that hosts the repository as the delegated administrator for CodeArtifact in the organization.
- C. Update the repository policy's Principal statement to include the ARN of the role that the CodeBuild project uses.
- D. Update the role that the CodeBuild project uses so that the role has sufficient permissions to use the CodeArtifact repository.
- E. Share the CodeArtifact repository with the organization by using AWS Resource Access Manager (AWS RAM).

Answer: A,D

Explanation:

Explanation

"AWS CodeArtifact operates in multiple Availability Zones and stores artifact data and metadata in Amazon S3 and Amazon DynamoDB. Your encrypted data is redundantly stored across multiple facilities and multiple devices in each facility, making it highly available and highly durable."

<https://aws.amazon.com/codeartifact/features/> With no internet connectivity, a gateway endpoint becomes necessary to access S3.

NEW QUESTION # 118

A company is refactoring applications to use AWS. The company identifies an internal web application that needs to make Amazon S3 API calls in a specific AWS account.

The company wants to use its existing identity provider (IdP) `auth.company.com` for authentication. The IdP supports only OpenID Connect (OIDC). A DevOps engineer needs to secure the web application's access to the AWS account.

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

- A. Create an IAM role that has a policy that allows the necessary S3 actions. Configure the role's trust policy to allow the OIDC IdP to assume the role if the `auth.company.com:aud` context key is `appid_from_idp`.
- B. Configure the web application to use the `GetFederationToken` API operation to retrieve temporary credentials Use the temporary credentials to make the S3 API calls.
- C. Configure AWS IAM Identity Center. Configure an IdP. Upload the IdP metadata from the existing IdP.
- D. Configure the web application to use the `AssumeRoleWithWebIdentity` API operation to retrieve temporary credentials. Use the temporary credentials to make the S3 API calls.
- E. Create an IAM role that has a policy that allows the necessary S3 actions. Configure the role's trust policy to allow the OIDC IdP to assume the role if the `sts.amazon.com:aud` context key is `appid from idp`.
- F. Create an IAM IdP by using the provider URL, audience, and signature from the existing IdP.

Answer: A,D,F

Explanation:

Step 1: Creating an Identity Provider in IAM You first need to configure AWS to trust the external identity provider (IdP), which in this case supports OpenID Connect (OIDC). The IdP will handle the authentication, and AWS will handle the authorization based on the IdP's token.

Action: Create an IAM Identity Provider (IdP) in AWS using the existing provider's URL, audience, and signature. This step is essential for establishing trust between AWS and the external IdP.

Why: This allows AWS to accept tokens from your external IdP (`auth.company.com`) for authentication.

Reference:

So, this corresponds to Option B: Create an IAM IdP by using the provider URL, audience, and signature from the existing IdP.

Step 2: Creating an IAM Role with Specific Permissions Next, you need to create an IAM role with a trust policy that allows the external IdP to assume it when certain conditions are met. Specifically, the trust policy needs to allow the role to be assumed based on the context key `auth.company.com:aud` (audience claim in the token).

Action: Create an IAM role that has the necessary permissions (e.g., Amazon S3 access). The role's trust policy should specify the OIDC IdP as the trusted entity and validate the audience claim (`auth.company.com:aud`), which comes from the token provided by the IdP.

Why: This step ensures that only the specified web application authenticated via OIDC can assume the IAM role to make API calls. This corresponds to Option D: Create an IAM role that has a policy that allows the necessary S3 actions. Configure the role's trust policy to allow the OIDC IdP to assume the role if the `auth.company.com:aud` context key is `appid_from_idp`.

Step 3: Using Temporary Credentials via `AssumeRoleWithWebIdentity` API To securely make Amazon S3 API calls, the web application will need temporary credentials. The web application can use the `AssumeRoleWithWebIdentity` API call to assume the

IAM role configured in the previous step and obtain temporary AWS credentials. These credentials can then be used to interact with Amazon S3.

Action: The web application must be configured to call the AssumeRoleWithWebIdentity API operation, passing the OIDC token from the IdP to obtain temporary credentials.

Why: This allows the web application to authenticate via the external IdP and then authorize access to AWS resources securely using short-lived credentials.

This corresponds to Option E: Configure the web application to use the AssumeRoleWithWebIdentity API operation to retrieve temporary credentials. Use the temporary credentials to make the S3 API calls.

Summary of Selected Answers:

B: Create an IAM IdP by using the provider URL, audience, and signature from the existing IdP.

D: Create an IAM role that has a policy that allows the necessary S3 actions. Configure the role's trust policy to allow the OIDC IdP to assume the role if the auth.company.com:aud context key is appid_from_idp.

E: Configure the web application to use the AssumeRoleWithWebIdentity API operation to retrieve temporary credentials. Use the temporary credentials to make the S3 API calls.

This setup enables the web application to use OpenID Connect (OIDC) for authentication and securely interact with Amazon S3 in a specific AWS account using short-lived credentials obtained through AWS Security Token Service (STS).

NEW QUESTION # 119

A company has many applications. Different teams in the company developed the applications by using multiple languages and frameworks. The applications run on premises and on different servers with different operating systems. Each team has its own release protocol and process. The company wants to reduce the complexity of the release and maintenance of these applications. The company is migrating its technology stacks, including these applications, to AWS. The company wants centralized control of source code, a consistent and automatic delivery pipeline, and as few maintenance tasks as possible on the underlying infrastructure. What should a DevOps engineer do to meet these requirements?

- A. Create one AWS CodeCommit repository for all applications. Put each application's code in a different branch. Merge the branches, and use AWS CodeBuild to build the applications. Use AWS CodeDeploy to deploy the applications to one centralized application server.
- B. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build the applications one at a time and to create one AMI for each server. Use AWS CloudFormation StackSets to automatically provision and decommission Amazon EC2 fleets by using these AMIs.
- C. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build the applications one at a time. Use AWS CodeDeploy to deploy the applications to one centralized application server.
- **D. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build one Docker image for each application in Amazon Elastic Container Registry (Amazon ECR). Use AWS CodeDeploy to deploy the applications to Amazon Elastic Container Service (Amazon ECS) on infrastructure that AWS Fargate manages.**

Answer: D

Explanation:

Explanation

because of "as few maintenance tasks as possible on the underlying infrastructure". Fargate does that better than "one centralized application server"

NEW QUESTION # 120

A company uses an organization in AWS Organizations that has all features enabled. The company uses AWS Backup in a primary account and uses an AWS Key Management Service (AWS KMS) key to encrypt the backups.

The company needs to automate a cross-account backup of the resources that AWS Backup backs up in the primary account. The company configures cross-account backup in the Organizations management account. The company creates a new AWS account in the organization and configures an AWS Backup backup vault in the new account. The company creates a KMS key in the new account to encrypt the backups. Finally, the company configures a new backup plan in the primary account. The destination for the new backup plan is the backup vault in the new account.

When the AWS Backup job in the primary account is invoked, the job creates backups in the primary account. However, the backups are not copied to the new account's backup vault.

Which combination of steps must the company take so that backups can be copied to the new account's backup vault? (Select TWO.)

- **A. Edit the backup vault access policy in the new account to allow access to the primary account.**

- B. Edit the key policy of the KMS key in the primary account to share the key with the new account.
- **C. Edit the key policy of the KMS key in the new account to share the key with the primary account.**
- D. Edit the backup vault access policy in the primary account to allow access to the KMS key in the new account.
- E. Edit the backup vault access policy in the primary account to allow access to the new account.

Answer: A,C

Explanation:

To enable cross-account backup, the company needs to grant permissions to both the backup vault and the KMS key in the destination account. The backup vault access policy in the destination account must allow the primary account to copy backups into the vault. The key policy of the KMS key in the destination account must allow the primary account to use the key to encrypt and decrypt the backups. These steps are described in the AWS documentation¹². Therefore, the correct answer is A and E.

Reference:

1: Creating backup copies across AWS accounts - AWS Backup

2: Using AWS Backup with AWS Organizations - AWS Backup

NEW QUESTION # 121

A company hosts applications in its AWS account. Each application logs to an individual Amazon CloudWatch log group. The company's CloudWatch costs for ingestion are increasing. A DevOps engineer needs to identify which applications are the source of the increased logging costs.

Which solution will meet these requirements?

- A. Use AWS CloudTrail to filter for CreateLogStream events for each application
- **B. Use AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage**
- C. Use CloudWatch Logs Insights to create a set of queries for the application log groups to identify the number of logs written for a period of time
- D. Use CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them.

Answer: B

Explanation:

The correct answer is C.

Option A is incorrect because using CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them is not a valid solution. CloudWatch metrics do not provide information about the size or volume of data being ingested by CloudWatch logs.

CloudWatch metrics only provide information about the number of events, bytes, and errors that occur within a log group or stream. Moreover, creating a custom expression with CloudWatch metrics would require using the `search_web` tool, which is not necessary for this use case.

Option B is incorrect because using CloudWatch Logs Insights to create a set of queries for the application log groups to identify the number of logs written for a period of time is not a valid solution. CloudWatch Logs Insights can help analyze and filter log events based on patterns and expressions, but it does not provide information about the cost or billing of CloudWatch logs. CloudWatch Logs Insights also charges based on the amount of data scanned by each query, which could increase the logging costs further.

Option C is correct because using AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage is a valid solution. AWS Cost Explorer is a tool that helps visualize, understand, and manage AWS costs and usage over time. AWS Cost Explorer can generate custom reports that show the breakdown of costs by service, region, account, tag, or any other dimension. AWS Cost Explorer can also filter and group costs by usage type, which can help identify the specific CloudWatch log groups that are the source of the increased logging costs.

Option D is incorrect because using AWS CloudTrail to filter for CreateLogStream events for each application is not a valid solution. AWS CloudTrail is a service that records API calls and account activity for AWS services, including CloudWatch logs. However, AWS CloudTrail does not provide information about the cost or billing of CloudWatch logs. Filtering for CreateLogStream events would only show when a new log stream was created within a log group, but not how much data was ingested or stored by that log stream.

References:

CloudWatch Metrics

CloudWatch Logs Insights

AWS Cost Explorer

AWS CloudTrail

DOWNLOAD the newest DumpExam DOP-C02 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1VITYzXONb8v7wTAuuiFOp4QNKXdFQ7_b