Latest DOP-C02 Dumps | Exam DOP-C02 Prep

In today's society, many people are busy every day and they think about changing their status of profession. They want to improve their competitiveness in the labor market, but they are worried that it is not easy to obtain the certification of DOP-C02. Our study tool can meet your needs. Once you use our DOP-C02 exam materials, you don't have to worry about consuming too much time, because high efficiency is our great advantage. You only need to spend 20 to 30 hours on practicing and consolidating of our DOP-C02 learning material, you will have a good result. After years of development practice, our DOP-C02 test torrent is absolutely the best. You will embrace a better future if you choose our DOP-C02 exam materials.

The DOP-C02 Certification Exam is a valuable certification for professionals who want to advance their careers in DevOps on the AWS platform. It is recognized by companies worldwide and is a prerequisite for several AWS partner programs. AWS Certified DevOps Engineer - Professional certification demonstrates the candidate's commitment to excellence and their ability to design, deploy, and manage highly available, fault-tolerant, and scalable AWS systems.

>> Latest DOP-C02 Dumps <<

Exam DOP-C02 Prep - Latest DOP-C02 Test Dumps

Preparation from reliable material is essential to get success in the real AWS Certified DevOps Engineer - Professional (DOP-C02) exam. One of the most crucial aspects of test preparation is relying on AWS Certified DevOps Engineer - Professional (DOP-C02) exam dumps. The authenticity of AWS Certified DevOps Engineer - Professional (DOP-C02) exam questions material plays a huge role in achieving a passing score. In the case of choosing, AWS Certified DevOps Engineer - Professional (DOP-C02) exam dumps outdated material, and one fails and loses resources. TestValid is committed to providing real DOP-C02 Questions, ensuring that applicants get success in a short time.

Amazon DOP-C02 (AWS Certified DevOps Engineer - Professional) Exam is a certification program offered by Amazon Web Services (AWS) to professionals who are interested in pursuing a career in the field of DevOps. AWS Certified DevOps Engineer -

Professional certification program is designed to validate the skills and expertise required to manage and deploy applications on the AWS platform using DevOps principles and practices.

Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q211-Q216):

NEW QUESTION # 211
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. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an OnFailure destination for the Lambda function
- B. Provision space in the /tmp folder of the Lambda function to give the function the ability to process large files from the S3 bucket
- C. Increase the memory of the Lambda function to give the function the ability to process large files from the S3 bucket.
- D. Create a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket

Answer: D

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 bucket1.
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 # 212
A company uses AWS CodePipeline pipelines to automate releases of its application A typical pipeline consists of three stages build, test, and deployment. The company has been using a separate AWS CodeBuild project to run scripts for each stage.
However, the company now wants to use AWS CodeDeploy to handle the deployment stage of the pipelines.
The company has packaged the application as an RPM package and must deploy the application to a fleet of Amazon EC2 instances. The EC2 instances are in an EC2 Auto Scaling group and are launched from a common AMI.

Which combination of steps should a DevOps engineer perform to meet these requirements? (Choose two.)

- A. Create an application in CodeDeploy. Configure an in-place deployment type. Specify the Auto Scaling group as the deployment target. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.
- B. Create a new version of the common AMI with the CodeDeploy agent installed. Create an AppSpec file that contains application deployment scripts and grants access to CodeDeploy.
- C. Create a new version of the common AMI with the CodeDeploy agent installed. Update the IAM role of the EC2 instances to allow access to CodeDeploy.
- D. Create an application in CodeDeploy. Configure an in-place deployment type. Specify the EC2 instances that are launched from the common AMI as the deployment target. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.
- E. Create an application in CodeDeploy. Configure an in-place deployment type. Specify the Auto Scaling group as the deployment target. Add a step to the CodePipeline pipeline to use EC2 Image Builder to create a new AMI. Configure CodeDeploy to deploy the newly created AMI.

Answer: A,C

Explanation:
Explanation
https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html

NEW QUESTION # 213
A company's application runs on Amazon EC2 instances. The application writes to a log file that records the username, date, time: and source IP address of the login. The log is published to a log group in Amazon CloudWatch Logs The company is performing a root cause analysis for an event that occurred on the previous day The company needs to know the number of logins for a specific user from the past 7 days Which solution will provide this information'?

- A. Create a CloudWatch Logs Insights query that uses an aggregation function to count the number of logins for the username over the past 7 days. Run the query against the log group
- B. Create a CloudWatch Logs metric filter on the log group Use a filter pattern that matches the username.
  Publish a CloudWatch metric that sums the number of logins over the past 7 days.
- C. Create a CloudWatch Logs subscription on the log group Use a filter pattern that matches the username Publish a CloudWatch metric that sums the number of logins over the past 7 days
- D. Create a CloudWatch dashboard. Add a number widget that has a filter pattern that counts the number of logins for the username over the past 7 days directly from the log group

Answer: A

Explanation:
To analyze and find the number of logins for a specific user from the past 7 days, a CloudWatch Logs Insights query is the most suitable solution. CloudWatch Logs Insights enables you to interactively search and analyze your log data in Amazon CloudWatch Logs. You can use the query language to perform queries that contain multiple commands, including aggregation functions, which can count the occurrences of logins for a specific username over a specified time period. This approach is more direct and efficient than creating a metric filter or subscription, which would require additional steps to publish and sum a metric. References: AWS Certified DevOps Engineer - Professional, CloudWatch Logs Insights query syntax, Tutorial: Run a query with an aggregation function, Add or remove a number widget from a CloudWatch dashboard.

NEW QUESTION # 214
A company has an application that runs on Amazon EC2 instances that are in an Auto Scaling group. When the application starts up. the application needs to process data from an Amazon S3 bucket before the application can start to serve requests.
The size of the data that is stored in the S3 bucket is growing. When the Auto Scaling group adds new instances, the application now takes several minutes to download and process the data before the application can serve requests. The company must reduce the time that elapses before new EC2 instances are ready to serve requests.
Which solution is the MOST cost-effective way to reduce the application startup time?

- A. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Running state. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
- B. Increase the maximum instance count of the Auto Scaling group. Configure an

autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.

- C. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Stopped state. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
- D. Increase the maximum instance count of the Auto Scaling group. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook and to place the new instance in the Standby state when the application is ready to serve requests.

Answer: C

Explanation:
Option A is the most cost-effective solution. By configuring a warm pool of EC2 instances in the Stopped state, the company can reduce the time it takes for new instances to be ready to serve requests. When the Auto Scaling group launches a new instance, it can attach the stopped EC2 instance from the warm pool. The instance can then be started up immediately, rather than having to wait for the data to be downloaded and processed. This reduces the overall startup time for the application.

NEW QUESTION # 215

A company's application teams use AWS CodeCommit repositories for their applications. The application teams have repositories in multiple AWS accounts. All accounts are in an organization in AWS Organizations.

Each application team uses AWS IAM Identity Center (AWS Single Sign-On) configured with an external IdP to assume a developer IAM role. The developer role allows the application teams to use Git to work with the code in the repositories.

A security audit reveals that the application teams can modify the main branch in any repository. A DevOps engineer must implement a solution that allows the application teams to modify the main branch of only the repositories that they manage.

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

- A. For each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.
- B. Create an approval rule template for each team in the Organizations management account. Associate the template with all the repositories. Add the developer role ARN as an approver.
- C. Create an approval rule template for each account. Associate the template with all repositories. Add the "aws:ResourceTag/access-team":"$;{aws:PrincipalTag/access-team}" condition to the approval rule template.
- D. Update the SAML assertion to pass the user's team name. Update the IAM role's trust policy to add an access-team session tag that has the team name.
- E. Attach an SCP to the accounts. Include the following statement: A computer code with text AI- generated content may be incorrect.
- F. Create an IAM permissions boundary in each account. Include the following statement: A computer code with black text AI-generated content may be incorrect.

Answer: A,D,F

Explanation:
Short Explanation: To meet the requirements, the DevOps engineer should update the SAML assertion to pass the user's team name, update the IAM role's trust policy to add an access-team session tag that has the team name, create an IAM permissions boundary in each account, and for each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.

Updating the SAML assertion to pass the user's team name allows the DevOps engineer to use IAM tags to identify which team a user belongs to. This can help enforce fine-grained access control based on the user's team membership1.

Updating the IAM role's trust policy to add an access-team session tag that has the team name allows the DevOps engineer to use IAM condition keys to restrict access based on the session tag value2. For example, the DevOps engineer can use the aws:PrincipalTag condition key to match the access-team tag of the user with the access-team tag of the repository3.

Creating an IAM permissions boundary in each account allows the DevOps engineer to set the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries4. For example, the DevOps engineer can use a permissions boundary policy to limit the actions that a user can perform on CodeCommit repositories based on their access-team tag5.

For each CodeCommit repository, adding an access-team tag that has the value set to the name of the associated team allows the DevOps engineer to use resource tags to identify which team manages a repository. This can help enforce fine-grained access control based on the resource tag value6.

The other options are incorrect because:

Creating an approval rule template for each team in the Organizations management account is not a valid option, as approval rule

templates are not supported by AWS Organizations. Approval rule templates are specific to CodeCommit and can only be associated with one or more repositories in the same AWS Region where they are created7.

Creating an approval rule template for each account is not a valid option, as approval rule templates are not designed to restrict access to modify branches. Approval rule templates are designed to require approvals from specified users or groups before merging pull requests8.

Attaching an SCP to the accounts is not a valid option, as SCPs are not designed to restrict access based on tags. SCPs are designed to restrict access based on service actions and resources across all users and roles in an organization's account9.

NEW QUESTION # 216

......

Exam DOP-C02 Prep: https://www.testvalid.com/DOP-C02-exam-collection.html