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.geekcert.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

When logging with Amazon CloudTrail, API call information for services with regional end points is ____.

A. captured and processed in the same region as to which the API call is made and delivered to the region associated with your Amazon S3 bucket

B. captured, processed, and delivered to the region associated with your Amazon S3 bucket

C. captured in the same region as to which the API call is made and processed and delivered to the region associated with your Amazon S3 bucket

D. captured in the region where the end point is located, processed in the region where the CloudTrail trail is configured, and delivered to the region associated with your Amazon S3 bucket

Correct Answer: A

When logging with Amazon CloudTrail, API call information for services with regional end points (EC2, RDS etc.) is captured and processed in the same region as to which the API call is made and delivered to the region associated with your Amazon S3 bucket. API call information for services with single end points (IAM, STS etc.) is captured in the region where the end point is located, processed in the region where the CloudTrail trail is configured, and delivered to the region associated with your Amazon S3 bucket.

Reference: https://aws.amazon.com/cloudtrail/faqs/

QUESTION 2

When building a Docker image, you are searching through a persistent data volume\\'s logs to provide parameters for the next build. You execute the following command. Which of the operations will cause a failure of the Docker RUNcommand? RUN cat ./data/log/*.error | grep service_status | grep ERROR

A. the first grep command

B. any one of them

C. the second grep command

D. the cat command

Correct Answer: C

Some RUN commands depend on the ability to pipe the output of one command into another, using the pipe character (|), as in the following example:

RUN wget -O - https://some.site | wc -l > /number

Docker executes these commands using the /bin/sh -c interpreter, which only evaluates the exit code of the last operation in the pipe to determine success. In the example above this build step succeeds and produces a new image so long as

the wc -lcommand succeeds, even if the wget command fails.

Reference:

https://docs.docker.com/engine/userguide/eng-image/dockerfile_best-practices/#run

QUESTION 3

If designing a single playbook to run across multiple Linux distributions that have distribution specific commands, what would be the best method to allow a successful run?

A. Enable fact gathering and use the `when\\' conditional to match the distribution to the task.

B. This is not possible, a separate playbook for each target Linux distribution is required.

C. Use `ignore_errors: true\\' in the tasks.

D. Use the `shell\\' module to write your own checks for each command that is ran.

Correct Answer: A

Ansible provides a method to only run a task when a condition is met using the `when\\' declarative. With gather facts enabled, the play has access to the distribution name of the Linux system, thus, tasks can be tailored to a specific distribution and ran only when the condition is met, e.g.: ` - when: ansible_os_family == "Debian"\\'.

Reference: http://docs.ansible.com/ansible/playbooks_conditionals.html

QUESTION 4

A company uses AWS CodeCommit for source code control. Developers apply their changes to various feature branches and create pull requests to move those changes to the main branch when the changes are ready for production.

The developers should not be able to push changes directly to the main branch. The company applied the AWSCodeCommitPowerUser managed policy to the developers\\' IAM role, and now these developers can push changes to the main branch directly on every repository in the AWS account.

What should the company do to restrict the developers\\' ability to push changes to the main branch directly?

A. Create an additional policy to include a Deny rule for the GitPush and PutFile actions. Include a restriction for the specific repositories in the policy statement with a condition that references the main branch.

B. emove the IAM policy, and add an AWSCodeCommitReadOnly managed policy. Add an Allow rule for the GitPush and PutFile actions for the specific repositories in the policy statement with a condition that references the main branch.

C. Modify the IAM policy. Include a Deny rule for the GitPush and PutFile actions for the specific repositories in the policy statement with a condition that references the main branch.

D. Create an additional policy to include an Allow rule for the GitPush and PutFile actions. Include a restriction for the specific repositories in the policy statement with a condition that references the feature branches.

Correct Answer: A

By default, the AWSCodeCommitPowerUser managed policy allows users to push changes to any branch in any repository in the AWS account. To restrict the developers\\' ability to push changes to the main branch directly, an additional

policy is needed that explicitly denies these actions for the main branch.

The Deny rule should be included in a policy statement that targets the specific repositories and includes a condition that references the main branch. The policy statement should look something like this:

{

"Effect": "Deny",

"Action": [

"codecommit:GitPush",

"codecommit:PutFile"

],

"Resource": "arn:aws:codecommit:::", "Condition": {

"StringEqualsIfExists": {

"codecommit:References": [

"refs/heads/main"

]

}

}

QUESTION 5

Amazon Inspector agent collects telemetry data during assessment run and sends this data to Amazon Inspector dedicated S3 bucket for analysis. How can you access telemetry data out of Amazon Inspector and how can you benefit from this data in securing your resources?

A. Telemetry data is kept in S3 and encrypted with a pre-assessment test key configured in KMS, as long as you have access to that key you can download and decrypt telemetry data.

B. Telemetry data is stored in Amazon Inspector dedicated S3 bucket that does NOT belong to your account, Amazon Inspector currently does NOT provide an API or an S3 bucket access mechanism to collected telemetry. Data is retained temporarily only to allow for assistance with support requests.

C. Telemetry data is saved on S3 bucket in your account, therefore telemetry data is accessible with proper permissions on that bucket.

D. Telemetry data is deleted immediately after assessment run, therefore data can NOT be accessed or analyzed by any other tools.

Correct Answer: B

The telemetry data stored in S3 is retained only to allow for assistance with support requests and is not used or aggregated by Amazon for any other purpose. After 30 days, telemetry data is permanently deleted per a standard Amazon Inspector-dedicated S3 bucket lifecycle policy. At present, Amazon Inspector does not provide an API or an S3

bucket access mechanism to collected telemetry.

Reference: https://docs.aws.amazon.com/inspector/latest/userguide/inspector_agents.html