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.pass4itsure.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

SATISFACTION GUARANTEED
100%
SATISFACTION GUARANTEED

QUESTION 1

A company\\'s development team uses AWS CloudFormation to deploy its application resources. The team must use CloudFormation for all changes to the environment. The team cannot use the AWS Management Console or the AWS CLI to make manual changes directly.

The team uses a developer IAM role to access the environment. The role is configured with the AdministratorAccess managed IAM policy. The company has created a new CloudFormationDeployment IAM role that has the following policy attached:

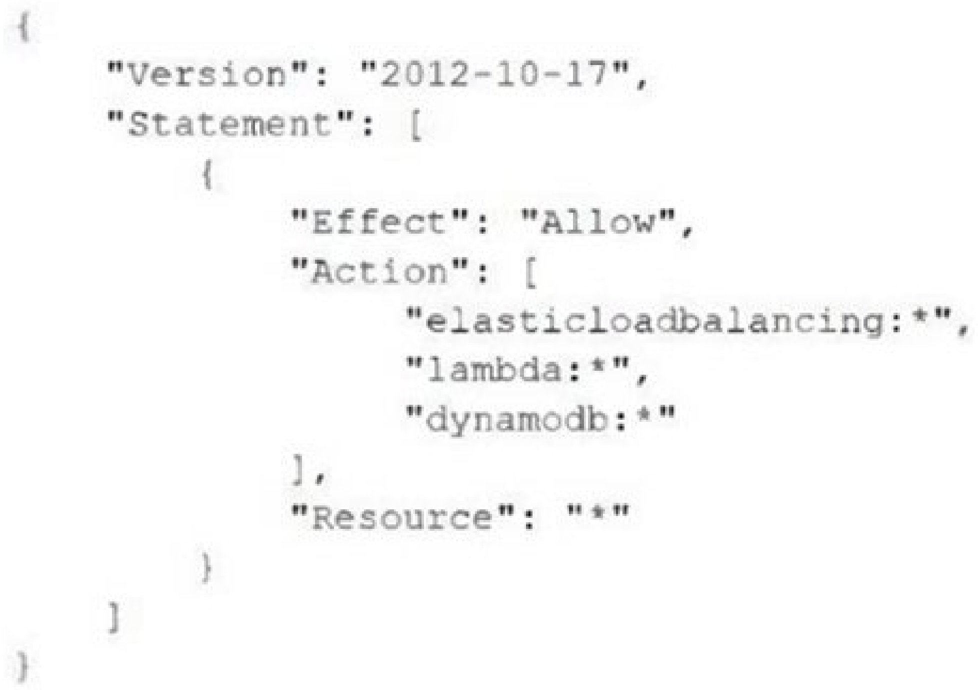

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "elasticloadbalancing:*",
                "lambda:*",
                "dynamodb:*"
            ],
            "Resource": "*"
        }
    ]
}
```

The company wants to ensure that only CloudFormation can use the new role. The development team cannot make any manual changes to the deployed resources.

Which combination of steps will meet these requirements? (Choose three.)

A. Remove the AdministratorAccess policy. Assign the ReadOnlyAccess managed IAM policy to the developer role. Instruct the developers to use the CloudFormationDeployment role as a CloudFormation service role when the developers deploy new stacks.

B. Update the trust policy of the CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDeployment role.

C. Configure the developer IAM role to be able to get and pass the CloudFormationDeployment role if iam:PassedToService equals . Configure the CloudFormationDeployment role to allow all cloudformation actions for all resources.

D. Update the trust policy of the CloudFormationDeployment role to allow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeRole action.

E. Remove the AdministratorAccess policy. Assign the ReadOnlyAccess managed IAM policy to the developer role. Instruct the developers to assume the CloudFormationDeployment role when the developers deploy new stacks.

F. Add an IAM policy to the CloudFormationDeployment role to allow cloudformation:* on all resources. Add a policy that allows the iam:PassRole action for the ARN of the CloudFormationDeployment role if iam:PassedToService equals cloudformation.amazonaws.com.

Correct Answer: ADF

QUESTION 2

A global company manages multiple AWS accounts by using AWS Control Tower. The company hosts internal applications and public applications.

Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations. One of the AWS Control Tower member accounts serves as a centralized DevOps account with CI/CD pipelines that application teams use to deploy applications to their respective target AWS accounts. An IAM role for deployment exists in the centralized DevOps account.

An application team is attempting to deploy its application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in an application AWS account. An IAM role for deployment exists in the application AWS account. The deployment is through an AWS CodeBuild project that is set up in the centralized DevOps account. The CodeBuild project uses an IAM service role for CodeBuild. The deployment is failing with an Unauthorized error during attempts to connect to the cross-account EKS cluster from CodeBuild.

Which solution will resolve this error?

A. Configure the application account\\'s deployment IAM role to have a trust relationship with the centralized DevOps account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the application account\\'s deployment IAM role to have the required access to the EKS cluster. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

B. Configure the centralized DevOps account\\'s deployment IAM role to have a trust relationship with the application account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the centralized DevOps account\\'s deployment IAM role to allow the required access to CodeBuild.

C. Configure the centralized DevOps account\\'s deployment IAM role to have a trust relationship with the application account. Configure the trust relationship to allow the sts:AssumeRoleWithSAML action. Configure the centralized DevOps account\\'s deployment IAM role to allow the required access to CodeBuild.

D. Configure the application account\\'s deployment IAM role to have a trust relationship with the AWS Control Tower management account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the application account\\'s deployment IAM role to have the required access to the EKS cluster. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

Correct Answer: A

In the source AWS account, the IAM role used by the CI/CD pipeline should have permissions to access the source code repository, build artifacts, and any other resources required for the build process. In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application. the IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.

QUESTION 3

A cloud team uses AWS Organizations and AWS IAM Identity Center (AWS Single Sign-On) to manage a company\\'s AWS accounts. The company recently established a research team. The research team requires the ability to fully manage the resources in its account. The research team must not be able to create IAM users.

The cloud team creates a Research Administrator permission set in IAM Identity Center for the research team. The permission set has the AdministratorAccess AWS managed policy attached. The cloud team must ensure that no one on the research team can create IAM users.

Which solution will meet these requirements?

A. Create an IAM policy that denies the iam:CreateUser action. Attach the IAM policy to the Research Administrator permission set.

B. Create an IAM policy that allows all actions except the iam:CreateUser action. Use the IAM policy to set the permissions boundary for the Research Administrator permission set.

C. Create an SCP that denies the iam:CreateUser action. Attach the SCP to the research team\\'s AWS account.

D. Create an AWS Lambda function that deletes IAM users. Create an Amazon EventBridge rule that detects the IAM CreateUser event. Configure the rule to invoke the Lambda function.

Correct Answer: C

QUESTION 4

A DevOps engineer has created an AWS CloudFormation template that deploys an application on Amazon EC2 instances. The EC2 instances run Amazon Linux. The application is deployed to the EC2 instances by using shell scripts that contain user data. The EC2 instances have an IAM instance profile that has an IAM role with the AmazonSSMManagedinstanceCore managed policy attached.

The DevOps engineer has modified the user data in the CloudFormation template to install a new version of the application. The engineer has also applied the stack update. However, the application was not updated on the running EC2 instances. The engineer needs to ensure that the changes to the application are installed on the running EC2 instances.

Which combination of steps will meet these requirements? (Choose two.)

A. Configure the user data content to use the Multipurpose Internet Mail Extensions (MIME) multipart format. Set the scripts-user parameter to always in the text/cloud-config section.

B. Refactor the user data commands to use the cfn-init helper script. Update the user data to install and configure the cfn-hup and cfn-init helper scripts to monitor and apply the metadata changes.

C. Configure an EC2 launch template for the EC2 instances. Create a new EC2 Auto Scaling group. Associate the Auto Scaling group with the EC2 launch template. Use the AutoScalingScheduledAction update policy for the Auto Scaling group.

D. Refactor the user data commands to use an AWS Systems Manager document (SSM document). Add an AWS CLI command in the user data to use Systems Manager Run Command to apply the SSM document to the EC2 instances.

E. Refactor the user data command to use an AWS Systems Manager document (SSM document). Use Systems Manager State Manager to create an association between the SSM document and the EC2 instances.

Correct Answer: BE

QUESTION 5

An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.

When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.

How should the company meet these requirements with the LEAST amount of application changes?

A. Use Amazon Redshift for the product catalog and Amazon DynamoDB tables for the customer information and purchases.

B. Use Amazon DynamoDB global tables for the product catalog and regional tables for the customer information and purchases.

C. Use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases.

D. Use Aurora for the product catalog and Amazon DynamoDB global tables for the customer information and purchases.

Correct Answer: C

QUESTION 6

A DevOps engineer needs to implement integration tests into an existing AWS CodePipeline CI/CD workflow for an Amazon Elastic Container Service (Amazon ECS) service. The CI/CD workflow retrieves new application code from an AWS

CodeCommit repository and builds a container image. The CI/CD workflow then uploads the container image to Amazon Elastic Container Registry (Amazon ECR) with a new image tag version.

The integration tests must ensure that new versions of the service endpoint are reachable and that various API methods return successful response data. The DevOps engineer has already created an ECS cluster to test the service.

Which combination of steps will meet these requirements with the LEAST management overhead? (Choose three.)

A. Add a deploy stage to the pipeline. Configure Amazon ECS as the action provider. Most Voted

B. Add a deploy stage to the pipeline. Configure AWS CodeDeploy as the action provider.

C. Add an appspec.yml file to the CodeCommit repository.

D. Update the image build pipeline stage to output an imagedefinitions.json file that references the new image tag. Most Voted

E. Create an AWS Lambda function that runs connectivity checks and API calls against the service. Integrate the Lambda function with CodePipeline by using a Lambda action stage. Most Voted

F. Write a script that runs integration tests against the service. Upload the script to an Amazon S3 bucket. Integrate the script in the S3 bucket with CodePipeline by using an S3 action stage.

Correct Answer: ADE

QUESTION 7

What option below is the geographic limit of an EC2 security group?

A. Security groups are global.

B. They are confined to Placement Groups.

C. They are confined to Regions.

D. They are confined to Availability Zones.

Correct Answer: C

A security group is tied to a region and can be assigned only to instances in the same region.

You can\\'t enable an instance to communicate with an instance outside its region using security group rules. Traffic from an instance in another region is seen as WAN bandwidth.

Reference:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html

QUESTION 8

Which is the proper syntax for referencing a variable\\'s value in an Ansible task?

A. ${variable_name}

B. { variable_name }

C. "{{ variable_name }}"

D. @variable_name

Correct Answer: C

We use the variable\\'s name to reference the variable which we encapsulate in curly brackets `{{ }}\\'; however, the YAML syntax dictates that a string beginning with a curly bracket denotes a dictionary value. To get around this, it is proper to wrap the variable declaration in quotes.

Reference: http://docs.ansible.com/ansible/playbooks_variables.html#hey-wait-a-yaml-gotcha

QUESTION 9

A company uses AWS Organizations to manage its AWS accounts. The organization root has an OU that is named Environments. The Environments OU has two child OUs that are named Development and Production, respectively.

The Environments OU and the child OUs have the default FullAWSAccess policy in place. A DevOps engineer plans to remove the FullAWSAccess policy from the Development OU and replace the policy with a policy that allows all actions on Amazon EC2 resources.

What will be the outcome of this policy replacement?

A. All users in the Development OU will be allowed all API actions on all resources.

B. All users in the Development OU will be allowed all API actions on EC2 resources. All other API actions will be denied.

C. All users in the Development OU will be denied all API actions on all resources.

D. All users in the Development OU will be denied all API actions on EC2 resources. All other API actions will be allowed.

Correct Answer: B

The key point is that "SCP inheritance works differently for Allow and Deny policies". Allowed policies are only inherited if the children don\\'t have any Allow policy. Once they have an allow policy, only actions defined in that policy will be allowed and no "Allow" policy will be inherited from the parent(s) OUs. What inherits is the implicit Deny policy which is a hidden policy sitting above all.

Check the tables in this link: https://aws.amazon.com/blogs/security/get-more-out-of-service-control-policies-in-a-multi-account-environment/

QUESTION 10

You have deployed an application to AWS which makes use of Autoscaling to launch new instances. You now want to change the instance type for the new instances. Which of the following is one of the action items to achieve this deployment?

A. Use Elastic Beanstalk to deploy the new application with the new instance type

B. Use Cloudformation to deploy the new application with the new instance type

C. Create a new launch configuration with the new instance type

D. Create new EC2 instances with the new instance type and attach it to the Autoscaling Group

Correct Answer: C

The ideal way is to create a new launch configuration, attach it to the existing Auto Scaling group, and terminate the running instances. Option A is invalid because Clastic beanstalk cannot launch new instances on demand. Since the current scenario requires Autoscaling, this is not the ideal option Option B is invalid because this will be a maintenance overhead, since you just have an Autoscaling Group. There is no need to create a whole Cloudformation template for this. Option D is invalid because Autoscaling Group will still launch CC2 instances with the older launch configuration.

QUESTION 11

A company deploys an application to Amazon EC2 instances. The application runs Amazon Linux 2 and uses AWS CodeDeploy. The application has the following file structure for its code repository:

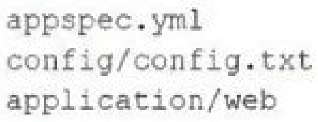

```
appspec.yml
config/config.txt
application/web
```

The appspec.yml file has the following contents in the files section:

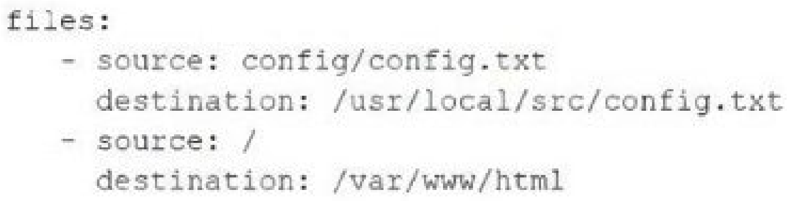

```
files:
    - source: config/config.txt
      destination: /usr/local/src/config.txt
    - source: /
      destination: /var/www/html
```

What will the result be for the deployment of the config.txt file?

A. The config.txt file will be deployed to only /var/www/html/config/config.txt.

B. The config.txt file will be deployed to /usr/local/src/config.txt and to /var/www/html/config/config.txt.

C. The config.txt file will be deployed to only /usr/local/src/config.txt.

D. The config.txt file will be deployed to /usr/local/src/config.txt and to /var/www/html/application/web/config.txt.

Correct Answer: B

QUESTION 12

When specifying multiple variable names and values for a playbook on the command line, which of the following is the correct syntax?

A. ansible-playbook playbook.yml -e `host="foo" pkg="bar"\\'

B. ansible-playbook playbook.yml -e `host: "foo", pkg: "bar"\\'

C. ansible-playbook playbook.yml -e `host="foo"\\' -e `pkg="bar"\\'

D. ansible-playbook playbook.yml --extra-vars "host=foo", "pkg=bar"

Correct Answer: A

Variables are passed in a single command line parameter, `-e\\' or `--extra-vars\\'. They are sent as a single string to the playbook and are space delimited. Because of the space delimeter, variable values must be encapsulated in quotes. Additionally, proper JSON or YAML can be passed, such as: `-e `{"key": "name", "array": ["value1", "value2"]}\\'.

Reference: http://docs.ansible.com/ansible/playbooks_variables.html#passing-variables-on-the-commandline

QUESTION 13

A company manages multiple AWS accounts by using AWS Organizations with OUS for the different business divisions, The company is updating their corporate network to use new IP address ranges. The company has 10 Amazon S3 buckets in different AWS accounts. The S3 buckets store reports for the different divisions. The S3 bucket configurations allow only private corporate network IP addresses to access the S3 buckets.

A DevOps engineer needs to change the range of IP addresses that have permission to access the contents of the S3 buckets The DevOps engineer also needs to revoke the permissions of two OUS in the company

Which solution will meet these requirements?

A. Create a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that demes access to the old range of IP addresses for all the S3 buckets. Set a permissions boundary for the OrganzauonAccountAccessRole role In the two OUS to deny access to the S3 buckets.

B. Create a new SCP that has a statement that allows only the new range of IP addresses to access the S3 buckets. Create another SCP that denies access to the S3 buckets. Attach the second SCP to the two OUS

C. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 buckets. Create a new SCP that denies access to the S3 buckets. Attach the SCP to the two OUs.

D. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 buckets. Set a permissions boundary for the OrganizationAccountAccessRole role in the two OUS to deny access to the S3 buckets.

Correct Answer: C

A comprehensive and detailed explanation is: Option A is incorrect because creating a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that denies access to the old range of IP addresses for all the S3 buckets, is not a valid solution. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. SCPs can only control the actions that can be performed by the principals in the organization, not the access to specific resources. Moreover, setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets. Option B is incorrect because creating a new SCP that has a statement that allows only the new range of IP addresses to access the S3 buckets is not a valid solution, for the same reason as option A. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. Creating another SCP that denies access to the S3 buckets and attaching it to the two OUs is also not a valid solution, as SCPs cannot specify the S3 buckets as resources either. Option C is correct because it meets both requirements of changing the range of IP addresses that have permission to access the contents of the S3 buckets and revoking the permissions of two OUs in the company. On all the S3 buckets, configuring resource-based policies that allow only the new range of IP addresses to access the S3 buckets is a valid way to update the IP address ranges, as resource-based policies can specify both resources and conditions. Creating a new SCP that denies access to the S3 buckets and attaching it to the two OUs is also a valid way to revoke the permissions of those OUs, as SCPs can deny actions such as s3:PutObject or s3:GetObject on any resource. Option D is incorrect because setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets. A permissions boundary is a policy that defines the maximum permissions that an IAM entity can have. However, it does not revoke any existing permissions that are granted by other policies. References: AWS Organizations S3 Bucket Policies Service Control Policies Permissions Boundaries

QUESTION 14

You are building a Docker image with the following Dockerfile. How many layers will the resulting image have?

FROM scratch CMD /app/hello.sh

A. 2

B. 4

C. 1

D. 3

Correct Answer: C

As described in the link: https://docs.docker.com/storage/storagedriver/, the DockerFile contain only one command which runs on the container layer. The container layer is a Read/Write layer. The CMD instruction specifies what command to run within the container, which only modifies the image\\'s metadata, which does not produce an image layer. So, there is only one layer i.e. the container layer.

QUESTION 15

An Amazon EC2 instance is running in a VPC and needs to download an object from a restricted Amazon S3 bucket. When the DevOps engineer tries to download the object, an AccessDenied error is received,

What are the possible causes tor this error? (Select TWO,)

A. The 53 bucket default encryption is enabled.

B. There is an error in the S3 bucket policy.

C. The object has been moved to S3 Glacier.

D. There is an error in the IAM role configuration.

E. S3 Versioning is enabled.

Correct Answer: BD

These are the possible causes for the AccessDenied error because they affect the permissions to access the S3 object from the EC2 instance. An S3 bucket policy is a resource-based policy that defines who can access the bucket and its objects, and what actions they can perform. An IAM role is an identity that can be assumed by an EC2 instance to grant it permissions to access AWS services and resources. If there is an error in the S3 bucket policy or the IAM role configuration, such as a missing or incorrect statement, condition, or principal, then the EC2 instance may not have the necessary permissions to download the object from the S3 bucket.
https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies.html
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

Latest DOP-C02 Dumps   DOP-C02 PDF Dumps   DOP-C02 Exam Questions