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

QUESTION 1

A DevOps engineer needs to apply a core set of security controls to an existing set of AWS accounts. The accounts are in an organization in AWS Organizations. Individual teams will administer individual accounts by using the AdministratorAccess AWS managed policy. For all accounts. AWS CloudTrail and AWS Config must be turned on in all available AWS Regions. Individual account administrators must not be able to edit or delete any of the baseline resources. However, individual account administrators must be able to edit or delete their own CloudTrail trails and AWS Config rules.

Which solution will meet these requirements in the MOST operationally efficient way?

A. Create an AWS CloudFormation template that defines the standard account resources. Deploy the template to all accounts from the organization\\'s management account by using CloudFormation StackSets. Set the stack policy to deny Update:Delete actions.

B. Enable AWS Control Tower. Enroll the existing accounts in AWS Control Tower. Grant the individual account administrators access to CloudTrail and AWS Config.

C. Designate an AWS Config management account. Create AWS Config recorders in all accounts by using AWS CloudFormation StackSets. Deploy AWS Config rules to the organization by using the AWS Config management account. Create a CloudTrail organization trail in the organization\\'s management account. Deny modification or deletion of the AWS Config recorders by using an SCP.

D. Create an AWS CloudFormation template that defines the standard account resources. Deploy the template to all accounts from the organization\\'s management account by using Cloud Formation StackSets Create an SCP that prevents updates or deletions to CloudTrail resources or AWS Config resources unless the principal is an administrator of the organization\\'s management account.

Correct Answer: D

QUESTION 2

Which one of the following is a restriction of AWS EBS Snapshots?

A. Snapshot restorations are restricted to the region in which the snapshots are created.

B. You cannot share unencrypted snapshots.

C. To share a snapshot with a user in other region the snapshot has to be created in that region first.

D. You cannot share a snapshot containing sensitive data such as an AWS Access Key ID or AWS Secret Access Key.

Correct Answer: C

Shapshots shared with other users are usable in full by the recipient, including but limited to the ability to base modified volumes and snapshots.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshotpermissions.html

QUESTION 3

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 4

A DevOps engineer needs to grant several external contractors access to a legacy application that runs on an Amazon Linux Amazon EC2 instance. The application server is available only in a private subnet. The contractors are not authorized for VPN access.

What should the DevOps engineer do to grant the contactors access to the application server?

A. Create an IAM user and SSH keys for each contractor. Add the public SSH key to the application server\\'s SSH authorized_keys file. Instruct the contractors to install the AWS CLI and AWS Systems Manager Session Manager plugin, update their AWS credentials files with their private keys, and use the aws ssm start-session command to gain access to the target application server instance ID.

B. Ask each contractor to securely send their SSH public key. Add this public key to the application server\\'s SSH authorized-keys file. Instruct the contractors to use their private key to connect to the application server through SSH.

C. Ask each contractor to securely send their SSH public key. Use EC2 pairs to import their key. Update the application server\\'s SSH authorized_keys file. Instruct the contractors to use their private key to connect to the application server through SSH.

D. Create an IAM user for each contractor with programmatic access. Add each user to an IAM group that has a policy that allows the ssm:StartSession action. Instruct the contractors to install the AWS CLI and AWS Systems Manager Session Manager plugin, update their AWS credentials files with their access keys, and use the aws ssm start-session to gain access to the target application server instance ID.

Correct Answer: B

QUESTION 5

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 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.

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 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.

D. 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.

E. 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.

Correct Answer: AD

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

QUESTION 6

A company wants to ensure that their EC2 instances are secure. They want to be notified if any new vulnerabilities are discovered on their instances and they also want an audit trail of all login activities on the instances.

Which solution will meet these requirements\\'?

A. Use AWS Systems Manager to detect vulnerabilities on the EC2 instances Install the Amazon Kinesis Agent to capture system logs and deliver them to Amazon S3.

B. Use AWS Systems Manager to detect vulnerabilities on the EC2 instances Install the Systems Manager Agent to capture system logs and view login activity in the CloudTrail console.

C. Configure Amazon CloudWatch to detect vulnerabilities on the EC2 instances Install the AWS Config daemon to capture system logs and view them in the AWS Config console.

D. Configure Amazon Inspector to detect vulnerabilities on the EC2 instances Install the Amazon CloudWatch Agent to capture system logs and record them via Amazon CloudWatch Logs.

Correct Answer: D

This solution will meet the requirements because it will use Amazon Inspector to scan the EC2 instances for any new vulnerabilities and generate findings that can be viewed in the Inspector console or sent as notifications via Amazon Simple Notification Service (SNS). It will also use the Amazon CloudWatch Agent to collect and send system logs from the EC2 instances to Amazon CloudWatch Logs, where they can be stored, searched, and analyzed. The system logs can provide an audit trail of all login activities on the instances, as well as other useful information such as performance metrics, errors, and events.

https://docs.aws.amazon.com/inspector/latest/user/what-is-inspector.html

QUESTION 7

A company is implementing an Amazon Elastic Container Service (Amazon ECS) cluster to run its workload. The company architecture will run multiple ECS services on the cluster. The architecture includes an Application Load Balancer on the front end and uses multiple target groups to route traffic.

A DevOps engineer must collect application and access logs. The DevOps engineer then needs to send the logs to an Amazon S3 bucket for near-real-time analysis.

Which combination of steps must the DevOps engineer take to meet these requirements? (Choose three.)

A. Download the Amazon CloudWatch Logs container instance from AWS. Configure this instance as a task. Update the application service definitions to include the logging task.

B. Install the Amazon CloudWatch Logs agent on the ECS instances. Change the logging driver in the ECS task definition to awslogs.

C. Use Amazon EventBridge to schedule an AWS Lambda function that will run every 60 seconds and will run the Amazon CloudWatch Logs create-export-task command. Then point the output to the logging S3 bucket.

D. Activate access logging on the ALB. Then point the ALB directly to the logging S3 bucket.

E. Activate access logging on the target groups that the ECS services use. Then send the logs directly to the logging S3 bucket.

F. Create an Amazon Kinesis Data Firehose delivery stream that has a destination of the logging S3 bucket. Then create an Amazon CloudWatch Logs subscription filter for Kinesis Data Firehose.

Correct Answer: BDF

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-logging-monitoring.html

QUESTION 8

Ansible supports running Playbook on the host directly or via SSH. How can Ansible be told to run its playbooks directly on the host?

A. Setting `connection: local\\' in the tasks that run locally.

B. Specifying `-type local\\' on the command line.

C. It does not need to be specified; it is the default.

D. Setting `connection: local\\' in the Playbook.

Correct Answer: D

Ansible can be told to run locally on the command line with the `-c\\' option or can be told via the `connection: local\\' declaration in the playbook. The default connection method is `remote\\'.

Reference: http://docs.ansible.com/ansible/intro_inventory.html#non-ssh-connection-types

QUESTION 9

A company uses Amazon EC2 instances to host applications for its customers. Recently, the company\\'s support team has received EC2 scheduled maintenance notifications regarding its EC2 instances.

The support team wants to automatically perform a restart of any EC2 instances with a scheduled maintenance event before the scheduled date.

Which solution will meet these requirements while requiring the MINIMUM amount of development effort?

A. Create an AWS Systems Manager maintenance window with a Systems Manager Automation tast that uses the RebootInstnaces EC2 API operation to restart the affected EC2 instances. Attach the EC2 instances to the maintenance window. Configure AWS Health to invoke the maintenance window whenever a scheduledChange event for Amazon EC2 is generated.

B. Create an Amazon CloudWatch alarm for the StatusCheckFailed metric of each EC2 instance. Configure the CloudWatch alarm to recover any affected EC2 instance.

C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that matches scheduledChange events for Amazon EC2 from AWS Health. Configure the rule to run the AWS-RestartEC2Instance AWS Systems Manager Automation runbook.

D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that matches scheduledChange events for

Amazon EC2 from AWS Health. Create an AWS Lambda function that uses the EC2 API to list all EC2 instances with scheduled events and then uses the RebootInstances EC2 API operation to restart the affected EC2 instances. Configure the EventBridge (CloudWatch Events) rule to invoke the Lambda function.

Correct Answer: A

QUESTION 10

The Ansible Inventory system allows many attributes to be defined within it. Which item below is not one of these?

A. Group variables

B. Host groups

C. Include vars

D. Children groups

Correct Answer: C

Ansible inventory files cannot reference other files for additional data. If this functionality is needed, it must be done in as a script to create a dynamic inventory.

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

QUESTION 11

A highly regulated company has a policy that DevOps engineers should not log in to their Amazon EC2 instances except in emergencies. It a DevOps engineer does log in the security team must be notified within 15 minutes of the occurrence.

Which solution will meet these requirements\\'?

A. Install the Amazon Inspector agent on each EC2 instance Subscribe to Amazon EventBridge notifications Invoke an AWS Lambda function to check if a message is about user logins If it is send a notification to the security team using

Amazon SNS.

B. Install the Amazon CloudWatch agent on each EC2 instance Configure the agent to push all logs to Amazon CloudWatch Logs and set up a CloudWatch metric filter that searches for user logins. If a login is found send a notification to the security team using Amazon SNS.

C. Set up AWS CloudTrail with Amazon CloudWatch Logs. Subscribe CloudWatch Logs to Amazon Kinesis Attach AWS Lambda to Kinesis to parse and determine if a log contains a user login If it does, send a notification to the security team using Amazon SNS.

D. Set up a script on each Amazon EC2 instance to push all logs to Amazon S3 Set up an S3 event to invoke an AWS Lambda function which invokes an Amazon Athena query to run. The Athena query checks tor logins and sends the output to the security team using Amazon SNS.

Correct Answer: B

https://aws.amazon.com/blogs/security/how-to-monitor-and-visualize-failed-ssh-access-attempts-to-amazon-ec2-linux-

instances/

QUESTION 12

A company has enabled all features for its organization in AWS Organizations. The organization contains 10 AWS accounts. The company has turned on AWS CloudTrail in all the accounts. The company expects the number of AWS accounts in the organization to increase to 500 during the next year. The company plans to use multiple OUs for these accounts.

The company has enabled AWS Config in each existing AWS account in the organization. A DevOps engineer must implement a solution that enables AWS Config automatically for all future AWS accounts that are created in the organization.

Which solution will meet this requirement?

A. In the organization\\'s management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call. Configure the rule to invoke an AWS Lambda function that enables trusted access to AWS Config for the organization.

B. In the organization\\'s management account, create an AWS CloudFormation stack set to enable AWS Config. Configure the stack set to deploy automatically when an account is created through Organizations.

C. In the organization\\'s management account, create an SCP that allows the appropriate AWS Config API calls to enable AWS Config. Apply the SCP to the root-level OU.

D. In the organization\\'s management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call. Configure the rule to invoke an AWS Systems Manager Automation runbook to enable AWS Config for the account.

Correct Answer: B

https://aws.amazon.com/about-aws/whats-new/2020/02/aws-cloudformation-stacksets-introduces-automatic-deployments-across-accounts-and-regions-through-aws-organizations/

QUESTION 13

A company has 20 service learns Each service team is responsible for its own microservice. Each service team uses a separate AWS account for its microservice and a VPC with the 192 168 0 0/22 CIDR block. The company manages the AWS accounts with AWS Organizations.

Each service team hosts its microservice on multiple Amazon EC2 instances behind an Application Load Balancer. The microservices communicate with each other across the public internet. The company\\'s security team has issued a new guideline that all communication between microservices must use HTTPS over private network connections and cannot traverse the public internet.

A DevOps engineer must implement a solution that fulfills these obligations and minimizes the number of changes for each service team

Which solution will meet these requirements?

A. Create a new AWS account in AWS Organizations Create a VPC in this account and use AWS Resource Access Manager to share the private subnets of this VPC with the organization Instruct the service teams to launch a new. Network Load Balancer (NLB) and EC2 instances that use the shared private subnets Use the NLB DNS names for

communication between microservices.

B. Create a Network Load Balancer (NLB) in each of the microservice VPCs Use AWS PrivateLink to create VPC endpoints in each AWS account for the NLBs Create subscriptions to each VPC endpoint in each of the other AWS accounts Use the VPC endpoint DNS names for communication between microservices.

C. Create a Network Load Balancer (NLB) in each of the microservice VPCs Create VPC peering connections between each of the microservice VPCs Update the route tables for each VPC to use the peering links Use the NLB DNS names for communication between microservices.

D. Create a new AWS account in AWS Organizations Create a transit gateway in this account and use AWS Resource Access Manager to share the transit gateway with the organization. In each of the microservice VPCs. create a transit gateway attachment to the shared transit gateway Update the route tables of each VPC to use the transit gateway Create a Network Load Balancer (NLB) in each of the microservice VPCs Use the NLB DNS names for communication between microservices.

Correct Answer: B

https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/ Private link is the best option because Transit Gateway doesn\\'t support overlapping CIDR ranges.

QUESTION 14

A company deploys updates to its Amazon API Gateway API several times a week by using an AWS CodePipeline pipeline. As part of the update process the company exports the JavaScript SDK for the API from the API. Gateway console and uploads the SDK to an Amazon S3 bucket

The company has configured an Amazon CloudFront distribution that uses the S3 bucket as an origin Web client then download the SDK by using the CloudFront distribution\\'s endpoint. A DevOps engineer needs to implement a solution to make the new SDK available automatically during new API deployments.

Which solution will meet these requirements?

A. Create a CodePipeline action immediately after the deployment stage of the API. Configure the action to invoke an AWS Lambda function. Configure the Lambda function to download the SDK from API Gateway, upload the SDK to the S3 bucket and create a CloudFront invalidation for the SDK path.

B. Create a CodePipeline action immediately after the deployment stage of the API Configure the action to use the CodePipelme integration with API. Gateway to export the SDK to Amazon S3 Create another action that uses the CodePipeline integration with Amazon S3 to invalidate the cache for the SDK path.

C. Create an Amazon EventBridge rule that reacts to UpdateStage events from aws apigateway Configure the rule to invoke an AWS Lambda function to download the SDK from API Gateway upload the SDK to the S3 bucket and call the CloudFront API to create an invalidation for the SDK path.

D. Create an Amazon EventBridge rule that reacts to Create. Deployment events from aws apigateway. Configure the rule to invoke an AWS Lambda function to download the SDK from API. Gateway upload the SDK to the S3 bucket and call the S3 API to invalidate the cache for the SDK path.

Correct Answer: A

This solution would allow the company to automate the process of updating the SDK and making it available to web clients. By adding a CodePipeline action immediately after the deployment stage of the API, the Lambda function will be invoked automatically each time the API is updated. The Lambda function should be able to download the new SDK from API Gateway, upload it to the S3 bucket and also create a CloudFront invalidation for the SDK path so that the

latest version of the SDK is available for the web clients. This is the most straight forward solution and it will meet the requirements.

QUESTION 15

Which of the following is an invalid variable name in Ansible?

A. host1st_ref

B. host-first-ref

C. Host1stRef

D. host_first_ref

Correct Answer: B

Variable names can contain letters, numbers and underscores and should always start with a letter. Invalid variable examples, `host first ref\\', `1st_host_ref\\'\\'.

Reference: http://docs.ansible.com/ansible/playbooks_variables.html#what-makes-a-valid-variable-name