



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 company deploys its corporate infrastructure on AWS across multiple AWS Regions and Availability Zones. The infrastructure is deployed on Amazon EC2 instances and connects with AWS IoT Greengrass devices. The company deploys additional resources on on-premises servers that are located in the corporate headquarters.

The company wants to reduce the overhead involved in maintaining and updating its resources. The company's DevOps team plans to use AWS Systems Manager to implement automated management and application of patches. The DevOps team confirms that Systems Manager is available in the Regions that the resources are deployed in. Systems Manager also is available in a Region near the corporate headquarters.

Which combination of steps must the DevOps team take to implement automated patch and configuration management across the company's EC2 instances IoT devices and on-premises infrastructure? (Select THREE.)

- A. Apply tags to all the EC2 instances, AWS IoT Greengrass devices, and on-premises servers. Use Systems Manager Session Manager to push patches to all the tagged devices.
- B. Use Systems Manager Run Command to schedule patching for the EC2 instances, AWS IoT Greengrass devices, and on-premises servers.
- C. Use Systems Manager Patch Manager to schedule patching for the EC2 instances, AWS IoT Greengrass devices, and on-premises servers as a Systems Manager maintenance window task.
- D. Configure Amazon EventBridge to monitor Systems Manager Patch Manager for updates to patch baselines. Associate Systems Manager Run Command with the event to initiate a patch action for all EC2 instances, AWS IoT Greengrass devices, and on-premises servers.
- E. Create an IAM instance profile for Systems Manager. Attach the instance profile to all the EC2 instances in the AWS account. For the AWS IoT Greengrass devices and on-premises servers, create an IAM service role for Systems Manager.
- F. Generate a managed-instance activation. Use the Activation Code and Activation ID to install Systems Manager Agent (SSM Agent) on each server in the on-premises environment. Update the AWS IoT Greengrass IAM token exchange role. Use the role to deploy SSM Agent on all the IoT devices.

Correct Answer: CEF

https://aws.amazon.com/blogs/mt/how-to-centrally-manage-aws-iot-greengrass-devices-using-aws-systems-manager/?force_isolation=true

QUESTION 2

A company wants to use a grid system for a proprietary enterprise memory data store on top of AWS. This system can run in multiple server nodes in any Linux-based distribution. The system must be able to reconfigure the entire cluster every time a node is added or removed. When adding or removing nodes, an `/etc./cluster/nodes` config file must be updated listing the IP addresses of the current node members of that cluster.

The company wants to automate the task of adding new nodes to a cluster.

What can a DevOps engineer do to meet these requirements?

- A. Use AWS OpsWorks Stacks to layer the server nodes of that cluster. Create a Chef recipe that populates the content of the `/etc./cluster/nodes` config file and restarts the service by using the current members of the layer. Assign that



recipe to the Configure lifecycle event.

B. Put the file nodes config in version control. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for the cluster nodes. When adding a new node to the cluster update the file with all tagged instances and make a commit in version control. Deploy the new file and restart the services.

C. Create an Amazon S3 bucket and upload a version of the `/etc./cluster/nodes` config file. Create a crontab script that will poll for that S3 file and download it frequently. Use a process manager such as Monit or system, to restart the cluster services when it detects that the new file was modified. When adding a node to the cluster edit the file's most recent members. Upload the new file to the S3 bucket.

D. Create a user data script that lists all members of the current security group of the cluster and automatically updates the `/etc/cluster/` nodes config. Tile whenever a new instance is added to the cluster.

Correct Answer: A

You can run custom recipes manually, but the best approach is usually to have AWS OpsWorks Stacks run them automatically. Every layer has a set of built-in recipes assigned each of five lifecycle events--Setup, Configure, Deploy, Undeploy, and Shutdown. Each time an event occurs for an instance, AWS OpsWorks Stacks runs the associated recipes for each of the instance's layers, which handle the corresponding tasks. For example, when an instance finishes booting, AWS OpsWorks Stacks triggers a Setup event. This event runs the associated layer's Setup recipes, which typically handle tasks such as installing and configuring packages

QUESTION 3

Using the AWS CLI, which command retrieves CloudTrail trail settings, including the status of the trail itself?

- A. `aws cloudtrail return-trails`
- B. `aws cloudtrail validate-settings`
- C. `aws cloudtrail get-settings`
- D. `aws cloudtrail describe-trails`

Correct Answer: D

You can retrieve trail settings and status using the `cloudtrail describe-trails` command. It will generate output similar to the example below.



```
{
  "trailList": [
    {
      "IncludeGlobalServiceEvents": false,
      "Name": "trailname",
      "S3KeyPrefix": "my-prefix",
      "TrailARN": "arn:aws:cloudtrail:us-east-2:123456789012:trail/my-trail",
      "LogFileValidationEnabled": true,
      "IsMultiRegionTrail": true,
      "HasCustomEventSelectors": false,
      "S3BucketName": " bucket"
      "SnsTopicName": " topic",
      "HomeRegion": "us-east-2"
    }
  ]
}
```

Reference: <http://docs.aws.amazon.com/awsccloudtrail/latest/userguide/cloudtrail-create-and-update-a-trailby-using-the-aws-cli.htm>

QUESTION 4

A DevOps Engineer has been asked to recommend a tool to deploy the components of a threetier web application. This application will use Amazon DynamoDB as a database Which deployment requires the LEAST amount of operational management?

- A. Use AWS CloudFormation to create a Classic Load Balancer and an Auto Scaling group. Use AWS OpsWorks to create the application and database resources Deploy application updates with OpsWorks using lifecycle events
- B. Use AWS OpsWorks to create a Classic Load Balancer, an Auto Scaling group application, and database resources Deploy application updates using OpsWorks lifecycle events
- C. Use AWS OpsWorks to create a Classic Load Balancer Auto Scaling and application resources Use AWS CloudFormation to create the database resources Deploy application updates using CloudFormation rolling updates
- D. Use AWS CloudFormation to create a Classic Load Balancer an Auto Scaling group and database resources Deploy application updates using CloudFormation rolling updates

Correct Answer: B

QUESTION 5

Which difference between core modules and extra modules is not correct?



- A. Extra modules may one day become core modules
- B. Core modules are supported by the Ansible team
- C. Core modules are shipped by default with Ansible
- D. Extra modules have no support

Correct Answer: D

While extra modules are not official modules and thus not supported by the Ansible team, they are indeed supported by their writers and the community.

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

[Latest DOP-C02 Dumps](#)

[DOP-C02 Study Guide](#)

[DOP-C02 Braindumps](#)