

DOP-C01^{Q&As}

AWS Certified DevOps Engineer - Professional (DOP-C01)

Pass Amazon DOP-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/aws-devops-engineer-professional.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 has a single Developer writing code for an automated deployment pipeline. The Developer is storing source code in an Amazon S3 bucket for each project. The company wants to add more Developers to the team but is concerned about code conflicts and lost work. The company also wants to build a test environment to deploy newer versions of code for testing and allow Developers to automatically deploy to both environments when code is changed in the repository.

What is the MOST efficient way to meet these requirements?

- A. Create an AWS CodeCommit repository for each project, use the main branch for production code, and create a testing branch for code deployed to testing. Use feature branches to develop new features and pull requests to merge code to testing and main branches.
- B. Create another S3 bucket for each project for testing code, and use an AWS Lambda function to promote code changes between testing and production buckets. Enable versioning on all buckets to prevent code conflicts.
- C. Create an AWS CodeCommit repository for each project, and use the main branch for production and test code with different deployment pipelines for each environment. Use feature branches to develop new features.
- D. Enable versioning and branching on each S3 bucket, use the main branch for production code, and create a testing branch for code deployed to testing. Have Developers use each branch for developing in each environment.

Correct Answer: A

QUESTION 2

An ecommerce company uses a large number of Amazon EBS backed Amazon EC2 instances. To decrease manual work across all the instances, a DevOps Engineer is tasked with automating restart actions when EC2 instance retirement events are scheduled.

How can this be accomplished?

- A. Create a scheduled Amazon CloudWatch Events rule to execute an AWS Systems Manager automation document that checks if any EC2 instances are scheduled for retirement once a week. If the instance is scheduled for retirement, the automation document will hibernate the instance.
- B. Enable EC2 Auto Recovery on all of the instances. Create an AWS Config rule to limit the recovery to occur during a maintenance window only.
- C. Reboot all EC2 instances during an approved maintenance window that is outside of standard business hours. Set up Amazon CloudWatch alarms to send a notification in case any instance is failing EC2 instance status checks.
- D. Set up an AWS Health Amazon CloudWatch Events rule to execute AWS Systems Manager automation documents that stop and start the EC2 instance when a retirement scheduled event occurs.

Correct Answer: D

Reference: <https://aws.amazon.com/blogs/mt/automate-remediation-actions-for-amazon-ec2-notifications-and-beyond-using-ec2-systems-manager-automation-and-aws-health/>

QUESTION 3

In DynamoDB, a secondary index is a data structure that contains a subset of attributes from a table, along with an alternate key to support _____ operations.

- A. None of the above
- B. Both
- C. Query
- D. Scan

Correct Answer: C

In DynamoDB, a secondary index is a data structure that contains a subset of attributes from a table, along with an alternate key to support Query operations.

QUESTION 4

A company wants to use a grid system for a proprietary enterprise in-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 systemd, 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` file whenever a new instance is added to the cluster

Correct Answer: A

QUESTION 5

A Development team is adding a new country to an e-commerce application. This addition requires that new application features be added to the shipping component of the application. The team has not decided if all new features should be

added, as some will take approximately six weeks to build. While the final decision on the shipping component features is being made, other team members are continuing to work on other features of the application.

Based on this situation, how should the application feature deployments be managed?

- A. Add the code updates as commits to the release branch. The team can delay the deployment until all features are ready.
- B. Add the code updates as commits to a feature branch. Merge the commits to a release branch as features are ready.
- C. Add the code updates as a single commit when a feature is ready. Tag this commit with "new-country."
- D. Create a new repository named "new-country". Commit all the code changes to the new repository.

Correct Answer: B

QUESTION 6

A company is using AWS CodePipeline to deploy an application. A recent policy change requires that a member of the company's security team sign off on any application changes before they are deployed into production. The approval should be recorded and retained.

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

- A. Configure CodePipeline with Amazon CloudWatch Logs to retain data.
- B. Configure CodePipeline to deliver action logs to Amazon S3.
- C. Create an AWS CloudTrail trail to deliver logs to Amazon S3.
- D. Create a custom CodePipeline action to invoke an AWS Lambda function for approval. Create a policy that gives the security team access to manage custom CodePipeline actions.
- E. Create a manual approval CodePipeline action before the deployment step. Create a policy that grants the security team access to approve manual approval stages.

Correct Answer: CE

QUESTION 7

You are running a Docker daemon on a Linux host and it becomes unresponsive. Which signal, when sent to a Docker process with the kill command, forces the full stack trace to be logged for debugging purposes?

- A. -TRACE
- B. -IOTRACE
- C. -SIGUSR1
- D. -KILLTRACE

Correct Answer: C

If the daemon is unresponsive, you can force a full stack trace to be logged by sending a SIGUSR1 signal to the daemon.

Linux:

```
$ sudo kill -SIGUSR1 $(pidof dockerd)
```

Windows Server:

Download docker-signal.

Run the executable with the flag --pid=.

Reference:

<https://docs.docker.com/engine/admin/#force-a-stack-trace-to-be-logged>

QUESTION 8

On which local address does the Docker DNS server listen?

- A. 127.0.0.1
- B. 127.0.0.111
- C. 127.0.0.254
- D. 127.0.0.11

Correct Answer: D

Note: If you need access to a host's localhost resolver, you must modify your DNS service on the host to listen on a non-localhost address that is reachable from within the container. Note: The DNS server is always at 127.0.0.11.

Reference: <https://docs.docker.com/engine/userguide/networking/configure-dns/>

QUESTION 9

A DevOps Engineer is architecting a continuous development strategy for a company's software as a service (SaaS) web application running on AWS. For application and security reasons, users subscribing to this application are distributed across multiple Application Load Balancers (ALBs), each of which has a dedicated Auto Scaling group and fleet of Amazon EC2 instances. The application does not require a build stage, and when it is committed to AWS CodeCommit, the application must trigger a simultaneous deployment to all ALBs, Auto Scaling groups, and EC2 fleets.

Which architecture will meet these requirements with the LEAST amount of configuration?

- A. Create a single AWS CodePipeline pipeline that deploys the application in parallel using unique AWS CodeDeploy applications and deployment groups created for each ALB-Auto Scaling group pair.
- B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.
- C. Create a single AWS CodePipeline pipeline that deploys the application in parallel using a single AWS CodeDeploy application and unique deployment group for each ALB-Auto Scaling group pair.
- D. Create an AWS CodePipeline pipeline for each ALB-Auto Scaling group pair that deploys the application using an AWS CodeDeploy application and deployment group created for the same ALB- Auto Scaling group pair.

Correct Answer: C

QUESTION 10

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

QUESTION 11

A DevOps Engineer is building a continuous deployment pipeline for a serverless application using AWS CodePipeline and AWS CodeBuild. The source, build, and test stages have been created with the deploy stage remaining. The company wants to reduce the risk of an unsuccessful deployment by deploying to a small percentage of customers and monitoring this deployment prior to a full release to all customers. How should the deploy stage be configured to meet these requirements?

- A. Use AWS CloudFormation to publish a new version on every stack update. Then set up a CodePipeline approval action for a Developer to test and approve the new version. Finally, use a CodePipeline invoke action to update an AWS Lambda function to use the production alias
- B. Use CodeBuild to use the AWS CLI to update the AWS Lambda function code, then publish a new version of the function and update the production alias to point to the new version of the function.
- C. Use AWS CloudFormation to define the serverless application and AWS CodeDeploy to deploy the AWS Lambda functions using DeploymentPreference: Canary10Percent15Minutes.
- D. Use AWS CloudFormation to publish a new version on every stack update. Use the RoutingConfig property of the AWS::Lambda::Alias resource to update the traffic routing during the stack update.

Correct Answer: C

QUESTION 12

Your application uses CloudFormation to orchestrate your application's resources. During your testing phase before the application went live, your Amazon RDS instance type was changed and caused the instance to be re-created, resulting in the loss of test data. How should you prevent this from occurring in the future?

- A. Within the AWS CloudFormation parameter with which users can select the Amazon RDS instance type, set AllowedValues to only contain the current instance type.
- B. Use an AWS CloudFormation stack policy to deny updates to the instance. Only allow UpdateStack permission to IAM principals that are denied SetStackPolicy.
- C. In the AWS CloudFormation template, set the AWS::RDS::DBInstance's DBInstanceClass property to be read-only.
- D. Subscribe to the AWS CloudFormation notification "BeforeResourceUpdate," and call CancelStackUpdate if the resource identified is the Amazon RDS instance.
- E. In the AWS CloudFormation template, set the DeletionPolicy of the AWS::RDS::DBInstance's DeletionPolicy property to "Retain."

Correct Answer: E

QUESTION 13

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 14

A company is implementing a well-architected design for its globally accessible API stack. The design needs to ensure both high reliability and fast response times for users located in North America and Europe.

The API stack contains the following three tiers:

- 1.

Amazon API Gateway

2.

AWS Lambda

3.

Amazon DynamoDB

Which solution will meet the requirements?

- A. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using health checks. Configure the APIs to forward requests to a Lambda function in that Region. Configure the Lambda functions to retrieve and update the data in a DynamoDB table in the same Region as the Lambda function.
- B. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using latency- based routing and health checks. Configure the APIs to forward requests to a Lambda function in that Region. Configure the Lambda functions to retrieve and update the data in a DynamoDB global table.
- C. Configure Amazon Route 53 to point to API Gateway in North America, create a disaster recovery API in Europe, and configure both APIs to forward requests to the Lambda functions in that Region. Retrieve the data from a DynamoDB global table. Deploy a Lambda function to check the North America API health every 5 minutes. In the event of a failure, update Route 53 to point to the disaster recovery API.
- D. Configure Amazon Route 53 to point to API Gateway API in North America using latency-based routing. Configure the API to forward requests to the Lambda function in the Region nearest to the user. Configure the Lambda function to retrieve and update the data in a DynamoDB table.

Correct Answer: B

QUESTION 15

A Developer is maintaining a fleet of 50 Amazon EC2 Linux servers. The servers are part of an Amazon EC2 Auto Scaling group, and also use Elastic Load Balancing for load balancing. Occasionally, some application servers are being

terminated after failing ELB HTTP health checks. The Developer would like to perform a root cause analysis on the issue, but before being able to access application logs, the server is terminated.

How can log collection be automated?

- A. Use Auto Scaling lifecycle hooks to put instances in a Pending:Wait state. Create an Amazon CloudWatch Alarm for EC2 Instance Terminate Successful and trigger an AWS Lambda function that executes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
- B. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create a Config rule for EC2 Instance-terminate Lifecycle Action and trigger a step function that executes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
- C. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create an Amazon CloudWatch subscription filter for EC2 Instance Terminate Successful and trigger a CloudWatch agent that executes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
- D. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create an Amazon CloudWatch Events

rule for EC2 Instance-terminate Lifecycle Action and trigger an AWS Lambda function that executes a SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.

Correct Answer: D

[DOP-C01 PDF Dumps](#)

[DOP-C01 VCE Dumps](#)

[DOP-C01 Practice Test](#)