



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.pass4itsure.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**

You are designing a system which needs, at minimum, 8 m4.large instances operating to service traffic. When designing a system for high availability in the us-east-1 region, which has 6 Availability Zones, your company needs to be able to handle the loss of a full availability zone. How should you distribute the servers, to save as much cost as possible, assuming all of the EC2 nodes are properly linked to an ELB?

Your VPC account can utilize us-east-1's AZs a through f, inclusive.

- A. 3 servers in each of AZs a through d, inclusive.
- B. 8 servers in each of AZs a and b.
- C. 2 servers in each of AZs a through e, inclusive.
- D. 4 servers in each of AZs a through c, inclusive.

Correct Answer: C

You need to design for N+1 redundancy on Availability Zones. $ZONE_COUNT = (REQUIRED_INSTANCES / INSTANCE_COUNT_PER_ZONE) + 1$. To minimize cost, spread the instances across as many possible zones as you can. By using a through e, you are allocating 5 zones. Using 2 instances, you have 10 total instances. If a single zone fails, you have 4 zones left, with 2 instances each, for a total of 8 instances. By spreading out as much as possible, you have increased cost by only 25% and significantly de-risked an availability zone failure.

Reference: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html>

QUESTION 2

For Amazon Inspector's integration with CloudTrail, what information is logged for List* and Describe* APIs?

- A. None. Amazon Inspector is an automated service and not monitored by CloudTrail.
- B. Both request and response information is logged.
- C. Only request information is logged.
- D. Request information is always logged. Response information is logged only for Completed assessment runs.

Correct Answer: C

For the Amazon Inspector integration with CloudTrail, for the List* and Describe* APIs, only the request information is logged.

Reference: <https://docs.aws.amazon.com/inspector/latest/userguide/logging-using-cloudtrail.html>

QUESTION 3

A DevOps Engineer must automate a weekly process of identifying unnecessary permissions on a per-user basis, across all users in an AWS account. This process should evaluate the permissions currently granted to each user by examining the user's attached IAM access policies compared to the permissions the user has actually used in the past



90 days. Any differences in the comparison would indicate that the user has more permissions than are required. A report of the deltas should be sent to the Information Security team for further review and IAM user access policy revisions, as required. Which solution is fully automated and will produce the MOST detailed deltas report?

- A. Create an AWS Lambda function that calls the IAM Access Advisor API to pull service permissions granted on a user-by-user basis for all users in the AWS account. Ensure that Access Advisor is configured with a tracking period of 90 days. Invoke the Lambda function using an Amazon CloudWatch Events rule on a weekly schedule. For each record, by user, by service, if the Access Advisor Last Accesses field indicates a day count instead of "Not accesses in the tracking period," this indicates a delta compared to what is in the user's currently attached access policies. After Lambda has iterated through all users in the AWS account, configure it to generate a report and send the report using Amazon SES.
- B. Configure an AWS CloudTrail trail that spans all AWS Regions and all read/write events, and point this trail to an Amazon S3 bucket. Create Amazon Athena table and specify the S3 bucket ARN in the CREATE TABLE query. Create an AWS Lambda function that accesses the Athena table using the SDK, which performs a SELECT, ensuring that the WHERE clause includes userIdentity, eventName, and eventTime. Compare the results against the user's currently attached IAM access policies to determine any deltas. Configure an Amazon CloudWatch Events schedule to automate this process to run once a week. Configure Amazon SES to send a consolidated report to the Information Security team.
- C. Configure VPC Flow Logs on all subnets across all VPCs in all regions to capture user traffic across the entire account. Ensure that all logs are being sent to a centralized Amazon S3 bucket, so all flow logs can be consolidated and aggregated. Create an AWS Lambda function that is triggered once a week by an Amazon CloudWatch Events schedule. Ensure that the Lambda function parses the flow log files for the following information: IAM user ID, subnet ID, VPC ID, Allow/Reject status per API call, and service name. Then have the function determine the deltas on a user-by-user basis. Configure the Lambda function to send the consolidated report using Amazon SES.
- D. Create an Amazon ES cluster and note its endpoint URL, which will be provided as an environment variable into a Lambda function. Configure an Amazon S3 event on a AWS CloudTrail trail destination S3 bucket and ensure that the event is configured to send to a Lambda function. Create the Lambda function to consume the events, parse the input from JSON, and transform it to an Amazon ES document format. POST the documents to the Amazon ES cluster's endpoint by way of the passed-in environment variable. Make sure that the proper indexing exists in Amazon ES and use Apache Lucene queries to parse the permissions on a user-by-user basis. Export the deltas into a report and have Amazon ES send the reports to the Information Security team using Amazon SES every week.

Correct Answer: B

QUESTION 4

What is true of the way that encryption works with EBS?

- A. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
- B. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
- C. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot always creates an encrypted volume.
- D. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot always creates an encrypted volume.

Correct Answer: C

Snapshots that are taken from encrypted volumes are automatically encrypted. Volumes that are created from



encrypted snapshots are also automatically encrypted. Your encrypted volumes and any associated snapshots always remain protected. For more information, see Amazon EBS Encryption. Reference:
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

QUESTION 5

When running a playbook on a remote target host you receive a Python error similar to "[Errno 13] Permission denied: `/home/nick/.ansible/tmp\'`. What would be the most likely cause of this problem?

- A. The user\'s home or `.ansible\' directory on the Ansible system is not writeable by the user running the play.
- B. The specified user does not exist on the remote system.
- C. The user running `ansible-playbook` must run it from their own home directory.
- D. The user\'s home or `.ansible\' directory on the Ansible remote host is not writeable by the user running the play

Correct Answer: D

Each task that Ansible runs calls a module. When Ansible uses modules, it copies the module to the remote target system. In the error above it attempted to copy it to the remote user\'s home directory and found that either the home directory or the `.ansible\' directory were not writeable and thus could not continue.

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

QUESTION 6

A company is using AWS for an application. The Development team must automate its deployments. The team has set up an AWS CodePipeline to deploy the application to Amazon EC2 instances by using AWS CodeDeploy after it has been built using the AWS CodeBuild service.

The team would like to add automated testing to the pipeline to confirm that the application is healthy before deploying it to the next stage of the pipeline using the same code. The team requires a manual approval action before the application is deployed, even if the test is successful. The testing and approval must be accomplished at the lowest costs, using the simplest management solution.

Which solution will meet these requirements?

- A. Add a manual approval action after the last deploy action of the pipeline. Use Amazon SNS to inform the team of the stage being triggered. Next, add a test action using CodeBuild to do the required tests. At the end of the pipeline, add a deploy action to deploy the application to the next stage.
- B. Add a test action after the last deploy action of the pipeline. Configure the action to use CodeBuild to perform the required tests. If these tests are successful, mark the action as successful. Add a manual approval action that uses Amazon SNS to notify the team, and add a deploy action to deploy the application to the next stage.
- C. Create a new pipeline that uses a source action that gets the code from the same repository as the first pipeline. Add a deploy action to deploy the code to a test environment. Use a test action using AWS Lambda to test the deployment. Add a manual approval action by using Amazon SNS to notify the team, and add a deploy action to deploy the application to the next stage.
- D. Add a test action after the last deployment action. Use a Jenkins server on Amazon EC2 to do the required tests and mark the action as successful if the tests pass. Create a manual approval action that uses Amazon SQS to notify the



team and add a deploy action to deploy the application to the next stage.

Correct Answer: B

QUESTION 7

When thinking of DynamoDB, what are true of Local Secondary Key properties?

- A. Either the partition key or the sort key can be different from the table, but not both.
- B. Only the sort key can be different from the table.
- C. The partition key and sort key can be different from the table.
- D. Only the partition key can be different from the table.

Correct Answer: B

Global secondary index ?an index with a partition key and a sort key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html>

QUESTION 8

A company wants to implement a CI/CD pipeline for an application that is deployed on AWS. The company also has a source-code analysis tool hosted on premises that checks for security flaws. The tool has not yet been migrated to AWS

and can be accessed only on-premises server. The company wants to run checks against the source code as part of the pipeline before the code is compiled. The checks take anywhere from minutes to an hour to complete.

How can a DevOps Engineer meet these requirements?

- A. Use AWS CodePipeline to create a pipeline. Add an action to the pipeline to invoke an AWS Lambda function after the source stage. Have the Lambda function invoke the source-code analysis tool on premises against the source input from CodePipeline. The function then waits for the execution to complete and places the output in a specified Amazon S3 location.
- B. Use AWS CodePipeline to create a pipeline, then create a custom action type. Create a job worker for the on-premises server that polls CodePipeline for job requests, initiates the tests, and returns the results. Configure the pipeline to invoke the custom action after the source stage.
- C. Use AWS CodePipeline to create a pipeline. Add a step after the source stage to make an HTTPS request to the on-premises hosted web service that invokes a test with the source code analysis tool. When the analysis is complete, the web service sends the results back by putting the results in an Amazon S3 output location provided by CodePipeline.
- D. Use AWS CodePipeline to create a pipeline. Create a shell script that copies the input source code to a location on premises. Invoke the source code analysis tool and return the results to CodePipeline. Invoke the shell script by adding a custom script action after the source stage.

Correct Answer: B



QUESTION 9

An Application team has three environments for their application: development, pre-production, and production. The team recently adopted AWS CodePipeline. However, the team has had several deployments of misconfigured or nonfunctional development code into the production environment, resulting in user disruption and downtime. The DevOps Engineer must review the pipeline and add steps to identify problems with the application before it is deployed.

What should the Engineer do to identify functional issues during the deployment process? (Choose two.)

- A. Use Amazon Inspector to add a test action to the pipeline. Use the Amazon Inspector Runtime Behavior Analysis Inspector rules package to check that the deployed code complies with company security standards before deploying it to production.
- B. Using AWS CodeBuild to add a test action to the pipeline to replicate common user activities and ensure that the results are as expected before progressing to production deployment.
- C. Create an AWS CodeDeploy action in the pipeline with a deployment configuration that automatically deploys the application code to a limited number of instances. The action then pauses the deployment so that the QA team can review the application functionality. When the review is complete, CodeDeploy resumes and deploys the application to the remaining production Amazon EC2 instances.
- D. After the deployment process is complete, run a testing activity on an Amazon EC2 instance in a different region that accesses the application to simulate user behavior. If unexpected results occur, the testing activity sends a warning to an Amazon SNS topic. Subscribe to the topic to get updates.
- E. Add an AWS CodeDeploy action in the pipeline to deploy the latest version of the development code to pre-production. Add a manual approval action in the pipeline so that the QA team can test and confirm the expected functionality. After the manual approval action, add a second CodeDeploy action that deploys the approved code to the production environment.

Correct Answer: BE

QUESTION 10

When logging with Amazon CloudTrail, API call information for services with single 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: D

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 11

A DevOps Engineer needs to design and implement a backup mechanism for Amazon EFS. The Engineer is given the following requirements:

1. The backup should run on schedule.
2. The backup should be stopped if the backup window expires.
3. The backup should be stopped if the backup completes before the backup window.
4. The backup logs should be retained for further analysis.
5. The design should support highly available and fault-tolerant paradigms.
6. Administrators should be notified with backup metadata.

Which design will meet these requirements?

- A. Use AWS Lambda with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activity. Run backup scripts on Amazon EC2 in an Auto Scaling group. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading backup logs to Amazon S3. Use Amazon SNS to notify administrators with backup activity metadata.
- B. Use Amazon SWF with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activity. Run backup scripts on Amazon EC2 in an Auto Scaling group. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading backup logs to Amazon Redshift. Use CloudWatch Alarms to notify administrators with backup activity metadata.
- C. Use AWS Data Pipeline with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activity. Run backup scripts on Amazon EC2 in a single Availability Zone. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading the backup logs to Amazon RDS. Use Amazon SNS to notify administrators with backup activity metadata.
- D. Use AWS CodePipeline with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activity. Run backup scripts on Amazon EC2 in a single Availability Zone. Use Auto Scaling lifecycle hooks and the SSM Run Command on Amazon EC2 for uploading backup logs to Amazon S3. Use Amazon SES to notify admins with backup activity metadata.

Correct Answer: A

QUESTION 12

A company has an application that has predictable peak traffic times. The company wants the application instances to scale up only during the peak times. The application stores state in Amazon DynamoDB. The application environment uses a standard Node.js application stack and custom Chef recipes stored in a private Git repository.

Which solution is MOST cost-effective and requires the LEAST amount of management overhead when performing rolling updates of the application environment?



- A. Create a custom AMI with the Node.js environment and application stack using Chef recipes. Use the AMI in an Auto Scaling group and set up scheduled scaling for the required times, then set up an Amazon EC2 IAM role that provides permission to access DynamoDB.
- B. Create a Docker file that uses the Chef recipes for the application environment based on an official Node.js Docker image. Create an Amazon ECS cluster and a service for the application environment, then create a task based on this Docker image. Use scheduled scaling to scale the containers at the appropriate times and attach a task-level IAM role that provides permission to access DynamoDB.
- C. Configure AWS OpsWorks stacks and use custom Chef cookbooks. Add the Git repository information where the custom recipes are stored, and add a layer in OpsWorks for the Node.js application server. Then configure the custom recipe to deploy the application in the deploy step. Configure time-based instances and attach an Amazon EC2 IAM role that provides permission to access DynamoDB.
- D. Configure AWS OpsWorks stacks and push the custom recipes to an Amazon S3 bucket and configure custom recipes to point to the S3 bucket. Then add an application layer type for a standard Node.js application server and configure the custom recipe to deploy the application in the deploy step from the S3 bucket. Configure time-based instances and attach an Amazon EC2 IAM role that provides permission to access DynamoDB.

Correct Answer: C

QUESTION 13

An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function using reserved concurrency. The Lambda function processes order messages from an Amazon SQS queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has Auto Scaling enabled for read and write capacity.

Which actions will diagnose and resolve the delay? (Choose two.)

- A. Check the ApproximateAgeOfOldestMessage metric for the SQS queue and increase the Lambda function concurrency limit.
- B. Check the ApproximateAgeOfOldestMessage metric for the SQS queue and configure a redrive policy on the SQS queue.
- C. Check the NumberOfMessagesSent metric for the SQS queue and increase the SQS queue visibility timeout.
- D. Check the ThrottledWriteRequests metric for the DynamoDB table and increase the maximum write capacity units for the table's Auto Scaling policy.
- E. Check the Throttles metric for the Lambda function and increase the Lambda function timeout.

Correct Answer: CE

QUESTION 14

In which Docker Swarm model does the swarm manager distribute a specific number of replica tasks among the nodes based upon the scale you set in the desired state?

- A. distributed services
- B. scaled services



C. replicated services

D. global services

Correct Answer: C

A service is the definition of the tasks to execute on the worker nodes. It is the central structure of the swarm system and the primary root of user interaction with the swarm. When you create a service, you specify which container image to use and which commands to execute inside running containers. In the replicated services model, the swarm manager distributes a specific number of replica tasks among the nodes based upon the scale you set in the desired state. For global services, the swarm runs one task for the service on every available node in the cluster. A task carries a Docker container and the commands to run inside the container. It is the atomic scheduling unit of swarm. Manager nodes assign tasks to worker nodes according to the number of replicas set in the service scale. Once a task is assigned to a node, it cannot move to another node. It can only run on the assigned node or fail.

Reference: <https://docs.docker.com/engine/swarm/key-concepts/#services-and-tasks>

QUESTION 15

You have a high security requirement for your AWS accounts. What is the most rapid and sophisticated setup you can use to react to AWS API calls to your account?

A. Subscription to AWS Config via an SNS Topic. Use a Lambda Function to perform in-flight analysis and reactivity to changes as they occur.

B. Global AWS CloudTrail setup delivering to S3 with an SNS subscription to the deliver notifications, pushing into a Lambda, which inserts records into an ELK stack for analysis.

C. Use a CloudWatch Rule ScheduleExpression to periodically analyze IAM credential logs. Push the deltas for events into an ELK stack and perform ad-hoc analysis there.

D. CloudWatch Events Rules which trigger based on all AWS API calls, submitting all events to an AWS Kinesis Stream for arbitrary downstream analysis.

Correct Answer: D

CloudWatch Events allow subscription to AWS API calls, and direction of these events into Kinesis Streams. This allows a unified, near real-time stream for all API calls, which can be analyzed with any tool (s) of your choosing downstream.

Reference:

http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/EventTypes.html#api_event_type

[DOP-C01 PDF Dumps](#)

[DOP-C01 VCE Dumps](#)

[DOP-C01 Practice Test](#)