

DUMPSBOSS.COM

AWS DevOps Engineer Professional (DOP-C01)

Amazon AWS AWS-DevOps-Engineer-Professional-DOP-C01

Version Demo

Total Demo Questions: 20

Total Premium Questions: 557

Buy Premium PDF

<https://dumpsboss.com>

support@dumpsboss.com

dumpsboss.com

QUESTION NO: 1

What storage driver does Docker generally recommend that you use if it is available?

- A. zfs
- B. btrfs
- C. aufs
- D. overlay

ANSWER: C**Explanation:**

After you have read the storage driver overview, the next step is to choose the best storage driver for your workloads. In making this decision, there are three high-level factors to consider: If multiple storage drivers are supported in your kernel, Docker has a prioritized list of which storage driver to use if no storage driver is explicitly configured, assuming that the prerequisites for that storage driver are met: If aufs is available, default to it, because it is the oldest storage driver. However, it is not universally available.

Reference:

<https://docs.docker.com/engine/userguide/storagedriver/selectadriver/>

QUESTION NO: 2

You have a large number of web servers in an Auto Scaling group behind a load balancer. On an hourly basis, you want to filter and process the logs to collect data on unique visitors, and then put that data in a durable data store in order to run reports. Web servers in the Auto Scaling group are constantly launching and terminating based on your scaling policies, but you do not want to lose any of the log data from these servers during a stop/termination initiated by a user or by Auto Scaling.

What two approaches will meet these requirements? (Choose two.)

- A.** Install an Amazon Cloudwatch Logs Agent on every web server during the bootstrap process. Create a CloudWatch log group and define Metric Filters to create custom metrics that track unique visitors from the streaming web server logs. Create a scheduled task on an Amazon EC2 instance that runs every hour to generate a new report based on the Cloudwatch custom metrics.
- B.** On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to Amazon Glacier. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminated. Use Amazon Data Pipeline to process the data in Amazon Glacier and run reports every hour.
- C.** On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to an Amazon S3 bucket. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminated. Use AWS Data Pipeline to move log data from the Amazon S3 bucket to Amazon Redshift in order to process and run reports every hour.

D. Install an AWS Data Pipeline Logs Agent on every web server during the bootstrap process. Create a log group object in AWS Data Pipeline, and define Metric Filters to move processed log data directly from the web servers to Amazon Redshift and run reports every hour.

ANSWER: A C

QUESTION NO: 3

A company is implementing an Amazon ECS cluster to run its workload. The company architecture will run multiple ECS services on the cluster, with an Application Load Balancer on the front end, using multiple target groups to route traffic. The Application Development team has been struggling to collect logs that must be collected and sent to an Amazon S3 bucket for near-real time analysis. What must the DevOps Engineer configure in the deployment to meet these requirements? (Choose three.)

- A.** Install the Amazon CloudWatch Logs logging agent on the ECS instances. Change the logging driver in the ECS task definition to 'awslogs'.
- B.** Download the Amazon CloudWatch Logs container instance from AWS and configure it as a task. Update the application service definitions to include the logging task.
- C.** Use Amazon CloudWatch Events to schedule an AWS Lambda function that will run every 60 seconds running the create-export -task CloudWatch Logs command, then point the output to the logging S3 bucket.
- D.** Enable access logging on the Application Load Balancer, then point it directly to the S3 logging bucket.
- E.** Enable access logging on the target groups that are used by the ECS services, then point it directly to the S3 logging bucket.
- F.** Create an Amazon Kinesis Data Firehose with a destination of the S3 logging bucket, then create an Amazon CloudWatch Logs subscription filter for Kinesis.

ANSWER: A D F

QUESTION NO: 4

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.

ANSWER: C E

QUESTION NO: 5

If you are trying to configure an AWS Elastic Beanstalk worker tier for easy debugging if there are problems finishing queue jobs, what should you configure?

- A. Configure Rolling Deployments
- B. Configure Enhanced Health Reporting
- C. Configure Blue-Green Deployments
- D. Configure a Dead Letter Queue

ANSWER: D

Explanation:

Elastic Beanstalk worker environments support Amazon Simple Queue Service (SQS) dead letter queues. A dead letter queue is a queue where other (source) queues can send messages that for some reason could not be successfully processed. A primary benefit of using a dead letter queue is the ability to sideline and isolate the unsuccessfully processed messages. You can then analyze any messages sent to the dead letter queue to try to determine why they were not successfully processed.

QUESTION NO: 6

You run a SIP-based telephony application that uses Amazon EC2 for its web tier and uses MySQL on Amazon RDS as its database. The application stores only the authentication profile data for its existing users in the database and therefore is read-intensive. Your monitoring system shows that your web instances and the database have high CPU utilization. Which of the following steps should you take in order to ensure the continual availability of your application? (Choose two.)

- A. Use a CloudFront RTMP download distribution with the application tier as the origin for the distribution.
- B. Set up an Auto Scaling group for the application tier and a policy that scales based on the Amazon EC2 CloudWatch CPU utilization metric.
- C. Vertically scale up the Amazon EC2 instances manually.
- D. Set up an Auto Scaling group for the application tier and a policy that scales based on the Amazon RDS CloudWatch CPU utilization metric.
- E. Switch to General Purpose (SSD) Storage from Provisioned IOPS Storage (PIOPS) for the Amazon RDS database.

F. Use multiple Amazon RDS read replicas.

ANSWER: B F

QUESTION NO: 7

A company wants to automatically re-create its infrastructure using AWS CloudFormation as part of the company's quality assurance (QA) pipeline. For each QA run, a new VPC must be created in a single account, resources must be deployed into the VPC, and tests must be run against this new infrastructure. The company policy states that all VPCs must be peered with a central management VPC to allow centralized logging. The company has existing CloudFormation templates to deploy its VPC and associated resources.

Which combination of steps will achieve the goal in a way that is automated and repeatable? (Choose two.)

- A.** Create an AWS Lambda function that is invoked by an Amazon CloudWatch Events rule when a `CreateVpcPeeringConnection` API call is made. The Lambda function should check the source of the peering request, accepts the request, and update the route tables for the management VPC to allow traffic to go over the peering connection.
- B.** In the CloudFormation template:
- Invoke a custom resource to generate unique VPC CIDR ranges for the VPC and subnets.
 - Create a peering connection to the management VPC.
 - Update route tables to allow traffic to the management VPC.
- C.** In the CloudFormation template:
- Use the `Fn::Cidr` function to allocate an unused CIDR range for the VPC and subnets.
 - Create a peering connection to the management VPC.
 - Update route tables to allow traffic to the management VPC.
- D.** Modify the CloudFormation template to include a mappings object that includes a list of /16 CIDR ranges for each account where the stack will be deployed.
- E.** Use CloudFormation StackSets to deploy the VPC and associated resources to multiple AWS accounts using a custom resource to allocate unique CIDR ranges. Create peering connections from each VPC to the central management VPC and accept those connections in the management VPC.

ANSWER: A B

QUESTION NO: 8

A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps Engineer is tasked with minimizing application response times and improving availability for users in both Regions.

Which combination of actions should be taken to address the latency issues? (Choose three.)

- A.** Create a new DynamoDB table in the new Region with cross-Region replication enabled.

- B.** Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.
- C.** Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group.
- D.** Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.
- E.** Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.
- F.** Convert the DynamoDB table to a global table.

ANSWER: C D F

QUESTION NO: 9

You run a 2000-engineer organization. You are about to begin using AWS at a large scale for the first time. You want to integrate with your existing identity management system running on Microsoft Active Directory, because your organization is a power-user of Active Directory. How should you manage your AWS identities in the most simple manner?

- A.** Use a large AWS Directory Service Simple AD.
- B.** Use a large AWS Directory Service AD Connector.
- C.** Use an Sync Domain running on AWS Directory Service.
- D.** Use an AWS Directory Sync Domain running on AWS Lambda

ANSWER: B

Explanation:

You must use AD Connector as a power-user of Microsoft Active Directory. Simple AD only works with a subset of AD functionality. Sync Domains do not exist; they are made up answers. AD Connector is a directory gateway that allows you to proxy directory requests to your on-premises Microsoft Active Directory, without caching any information in the cloud. AD Connector comes in 2 sizes; small and large.

A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector is designed for larger organizations of up to 5,000 users.

Reference:

<https://aws.amazon.com/directoryservice/details/>

QUESTION NO: 10

You are doing a load testing exercise on your application hosted on AWS. While testing your Amazon RDS MySQL DB instance, you notice that when you hit 100% CPU utilization on it, your application becomes non-responsive. Your application is read-heavy.

What are methods to scale your data tier to meet the application's needs? (Choose three.)

- A.** Add Amazon RDS DB read replicas, and have your application direct read queries to them.
- B.** Add your Amazon RDS DB instance to an Auto Scaling group and configure your CloudWatch metric based on CPU utilization.
- C.** Use an Amazon SQS queue to throttle data going to the Amazon RDS DB instance.
- D.** Use ElastiCache in front of your Amazon RDS DB to cache common queries.
- E.** Shard your data set among multiple Amazon RDS DB instances.
- F.** Enable Multi-AZ for your Amazon RDS DB instance.

ANSWER: A D E

QUESTION NO: 11

A company is using AWS to deploy an application. The development team must automate the deployments. The team has created an AWS CodePipeline pipeline to deploy the application to Amazon EC2 instances using AWS CodeDeploy after it has been built using AWS CodeBuild.

The team wants to add automated testing to the pipeline to confirm that the application is healthy before deploying the code to the EC2 instances. The team also requires a manual approval action before the application is deployed, even if the tests are successful. The testing and approval must be accomplished at the lowest costs, using the simplest management solution.

Which solution will meet these requirements?

- A.** Create a manual approval action after the build action of the pipeline. Use Amazon SNS to inform the team of the stage being triggered. Next, add a test action using CodeBuild to perform the required tests. At the end of the pipeline, add a deploy action to deploy the application to the next stage.
- B.** Create a test action after the CodeBuild build 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.** Create a test action after the build action. Use a Jenkins server on Amazon EC2 to perform 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.

ANSWER: B

QUESTION NO: 12

A DevOps engineer has been tasked with ensuring that all Amazon S3 buckets, except for those with the word "public" in the name, allow access only to authorized users utilizing S3 bucket policies. The security team wants to be notified when a bucket is created without the proper policy and for the policy to be automatically updated.

Which solutions will meet these requirements?

- A.** Create a custom AWS Config rule that will trigger an AWS Lambda function when an S3 bucket is created or updated. Use the Lambda function to look for S3 buckets that should be private, but that do not have a bucket policy that enforces privacy. When such a bucket is found, invoke a remediation action and use Amazon SNS to notify the security team.
- B.** Create an Amazon EventBridge (Amazon CloudWatch Events) rule that triggers when an S3 bucket is created. Use an AWS Lambda function to determine whether the bucket should be private. If the bucket should be private, update the PublicAccessBlock configuration. Configure a second EventBridge (CloudWatch Events) rule to notify the security team using Amazon SNS when PutBucketPolicy is called.
- C.** Create an Amazon S3 event notification that triggers when an S3 bucket is created that does not have the word "public" in the name. Define an AWS Lambda function as a target for this notification and use the function to apply a new default policy to the S3 bucket. Create an additional notification with the same filter and use Amazon SNS to send an email to the security team.
- D.** Create an Amazon EventBridge (Amazon CloudWatch Events) rule that triggers when a new object is created in a bucket that does not have the word "public" in the name. Target and use an AWS Lambda function to update the PublicAccessBlock configuration. Create an additional notification with the same filter and use Amazon SNS to send an email to the security team.

ANSWER: D

QUESTION NO: 13

To monitor API calls against our AWS account by different users and entities, we can use _____ to create a history of calls in bulk for later review, and use _____ for reacting to AWS API calls in real-time.

- A.** AWS Config; AWS Inspector
- B.** AWS CloudTrail; AWS Config
- C.** AWS CloudTrail; CloudWatch Events
- D.** AWS Config; AWS Lambda

ANSWER: C

Explanation:

CloudTrail is a batch API call collection service, CloudWatch Events enables real-time monitoring of calls through the Rules object interface.

Reference: <https://aws.amazon.com/whitepapers/security-at-scale-governance-in-aws/>

QUESTION NO: 14

After a recent audit, a company decided to implement a new disaster recovery strategy for its Amazon S3 data and its MySQL database running on Amazon EC2. Management wants the ability to recover to a secondary AWS Region with an RPO under 5 seconds and an RTO under 1 minute.

Which actions will meet the requirements while MINIMIZING operational overhead? (Choose two.)

- A. Modify the application to write to both Regions at the same time when uploading objects to Amazon S3.
- B. Migrate the database to an Amazon Aurora multi-master in the primary and secondary Regions.
- C. Migrate the database to Amazon RDS with a read replica in the secondary Region.
- D. Migrate to Amazon Aurora Global Database.
- E. Set up S3 cross-Region replication with a replication SLA for the S3 buckets where objects are being put.

ANSWER: C E

Explanation:

Reference: <https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategy-with-amazon-rds/>

QUESTION NO: 15

A DevOps engineer is deploying a new version of a company's application in an AWS CodeDeploy deployment group associated with its Amazon EC2 instances. After some time, the deployment fails. The engineer realizes that all the events associated with the specific deployment ID are in a Skipped status, and code was not deployed in the instances associated with the deployment group.

What are valid reasons for this failure? (Choose two.)

- A. The networking configuration does not allow the EC2 instances to reach the internet via a NAT gateway or internet gateway, and the CodeDeploy endpoint cannot be reached.
- B. The IAM user who triggered the application deployment does not have permission to interact with the CodeDeploy endpoint.
- C. The target EC2 instances were not properly registered with the CodeDeploy endpoint.
- D. An instance profile with proper permissions was not attached to the target EC2 instances.
- E. The appspec.yml file was not included in the application revision.

ANSWER: B C

QUESTION NO: 16

A company uses AWS Organizations to manage multiple accounts. Information security policies require that all unencrypted Amazon EBS volumes be marked as non-compliant. A DevOps engineer needs to automatically deploy the solution and ensure that this compliance check is always present.

With solution will accomplish this?

- A.** Create an AWS CloudFormation template that defines an AWS Inspector rule to check whether EBS encryption is enabled. Save the template to an Amazon S3 bucket that has been shared with all accounts within the company. Update the account creation script pointing to the CloudFormation template in Amazon S3.
- B.** Create an AWS Config organizational rule to check whether EBS encryption is enabled and deploy the rule using the AWS CLI. Create and apply an SCP to prohibit stopping and deleting AWS Config across the organization.
- C.** Create an SCP in Organizations. Set the policy to prevent the launch of Amazon EC2 instances without encryption on the EBS volumes using a conditional expression. Apply the SCP to all AWS accounts. Use Amazon Athena to analyze the AWS CloudTrail output, looking for events that deny an ec2:RunInstances action.
- D.** Deploy an IAM role to all accounts from a single trusted account. Build a pipeline with AWS CodePipeline with a stage in AWS Lambda to assume the IAM role, and list all EBS volumes in the account. Publish a report to Amazon S3.

ANSWER: A

QUESTION NO: 17

What is the expected behavior if Ansible is called with 'ansible-playbook -i localhost playbook.yml'?

- A.** Ansible will attempt to read the inventory file named 'localhost'
- B.** Ansible will run the plays locally.
- C.** Ansible will run the playbook on the host named 'localhost'
- D.** Ansible won't run, this is invalid command line syntax

ANSWER: A

Explanation:

Ansible expects an inventory filename with the '-i' option, regardless if it is a valid hostname. For this to execute on the host 'localhost' resolves to, a comma needs to be appended to the end.

Reference:

http://docs.ansible.com/ansible/intro_inventory.html#inventory

QUESTION NO: 18

A DevOps Engineer is deploying an Amazon API Gateway API with an AWS Lambda function providing the backend functionality. The Engineer needs to record the source IP address and response status of every API call.

Which combination of actions should the DevOps Engineer take to implement this functionality? (Choose three.)

- A.** Configure AWS X-Ray to enable access logging for the API Gateway requests.

- B. Configure the API Gateway stage to enable access logging and choose a logging format.
- C. Create a new Amazon CloudWatch Logs log group or choose an existing log group to store the logs.
- D. Grant API Gateway permission to read and write logs to Amazon CloudWatch through an IAM role.
- E. Create a new Amazon S3 bucket or choose an existing S3 bucket to store the logs.
- F. Configure API Gateway to stream its log data to Amazon Kinesis.

ANSWER: B C D

QUESTION NO: 19

To override an allow in an IAM policy, you set the Effect element to _____.

- A. Block
- B. Stop
- C. Deny
- D. Allow

ANSWER: C

Explanation:

By default, access to resources is denied. To allow access to a resource, you must set the Effect element to Allow. To override an allow (for example, to override an allow that is otherwise in force), you set the Effect element to Deny.

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

QUESTION NO: 20

When a user is detaching an EBS volume from a running instance and attaching it to a new instance, which of the below mentioned options should be followed to avoid file system damage?

- A. Unmount the volume first
- B. Stop all the I/O of the volume before processing
- C. Take a snapshot of the volume before detaching
- D. Force Detach the volume to ensure that all the data stays intact

ANSWER: A

Explanation:

When a user is trying to detach an EBS volume, the user can either terminate the instance or explicitly remove the volume. It is a recommended practice to unmount the volume first to avoid any file system damage.

DUMPSBOSS.COM