

# AWS Certified DevOps Engineer - Professional

1. A company wants to migrate its content sharing web application hosted on Amazon EC2 to a serverless architecture. The company currently deploys changes to its application by creating a new Auto Scaling group of EC2 instances and a new Elastic Load Balancer, and then shifting the traffic away using an Amazon Route

53 weighted routing policy.

For its new serverless application, the company is planning to use Amazon API Gateway and AWS Lambda. The company will need to update its deployment processes to work with the new application. It will also need to retain the ability to test new features on a small number of users before rolling the features out to the entire user base.

Which deployment strategy will meet these requirements?

A. Use AWS CDK to deploy API Gateway and Lambda functions. When code needs to be changed, update the AWS CloudFormation stack and deploy the new version of the APIs and Lambda functions. Use a Route 53 failover routing policy for the canary release strategy.

B. Use AWS CloudFormation to deploy API Gateway and Lambda functions using Lambda function versions. When code needs to be changed, update the CloudFormation stack with the new Lambda code and update the API versions using a canary release strategy. Promote the new version when testing is complete.

C. Use AWS Elastic Beanstalk to deploy API Gateway and Lambda functions. When code needs to be changed, deploy a new version of the API and Lambda functions. Shift traffic gradually using an Elastic Beanstalk blue/green deployment.

D. Use AWS OpsWorks to deploy API Gateway in the service layer and Lambda functions in a custom layer. When code needs to be changed, use OpsWorks to perform a blue/green deployment and shift traffic gradually.

**Answer(s): B**

---

2. 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(s):** C D F

---

3. A DevOps engineer used an AWS CloudFormation custom resource to set up AD Connector. The AWS Lambda function executed and created AD Connector, but CloudFormation is not transitioning from `CREATE_IN_PROGRESS` to `CREATE_COMPLETE`. Which action should the engineer take to resolve this issue?

- A. Ensure the Lambda function code has exited successfully.
- B. Ensure the Lambda function code returns a response to the pre-signed URL.
- C. Ensure the Lambda function IAM role has `cloudformation:UpdateStack` permissions for the stack ARN.
- D. Ensure the Lambda function IAM role has `ds:ConnectDirectory` permissions for the AWS account.

**Answer(s):** B

---

4. A company plans to stop using Amazon EC2 key pairs for SSH access, and instead plans to use AWS Systems Manager Session Manager. To further enhance security, access to Session Manager must take place over a private network only.

Which combinations of actions will accomplish this? (Choose two.)

- A. Allow inbound access to TCP port 22 in all associated EC2 security groups from the VPC CIDR range.
- B. Attach an IAM policy with the necessary Systems Manager permissions to the existing IAM instance profile.
- C. Create a VPC endpoint for Systems Manager in the desired Region.
- D. Deploy a new EC2 instance that will act as a bastion host to the rest of the EC2 instance fleet.
- E. Remove any default routes in the associated route tables.

**Answer(s):** B C

---

5. A company runs an application with an Amazon EC2 and on-premises configuration. A DevOps Engineer needs to standardize patching across both environments. Company policy dictates that patching only happens during non-business hours.

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

- A. Add the physical machines into AWS Systems Manager using Systems Manager Hybrid Activations.
- B. Attach an IAM role to the EC2 instances, allowing them to be managed by AWS Systems Manager.
- C. Create IAM access keys for the on-premises machines to interact with AWS Systems Manager.
- D. Execute an AWS Systems Manager Automation document to patch the systems every hour.
- E. Use Amazon CloudWatch Events scheduled events to schedule a patch window.
- F. Use AWS Systems Manager Maintenance Windows to schedule a patch window.

**Answer(s):** A B F

---

6. A company has many applications. Different teams in the company developed the applications by using multiple languages and frameworks. The applications run on premises and on different servers with different operating systems. Each team has its own release protocol and process. The company wants to reduce the complexity of the release and maintenance of these applications. The company is migrating its technology stacks, including these applications, to AWS. The company wants centralized control of source code, a consistent and automatic delivery pipeline, and as few maintenance tasks as possible on the underlying infrastructure. What should a DevOps engineer do to meet these requirements?

A. Create one AWS CodeCommit repository for all applications. Put each application's code in different branch. Merge the branches, and use AWS CodeBuild to build the applications. Use AWS CodeDeploy to deploy the applications to one centralized application server.

B. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build the applications one at a time. Use AWS CodeDeploy to deploy the applications to one centralized application server.

C. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build the applications one at a time to create one AMI for each server. Use AWS CloudFormation StackSets to automatically provision and decommission Amazon EC2 instances by using these AMIs.

D. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build one Docker image for each application in Amazon Elastic Container Registry (Amazon ECR). Use AWS CodeDeploy to deploy the applications to Amazon Elastic Container Service (Amazon ECS) on infrastructure that AWS Fargate manages.

**Answer(s):** D

---

7. A DevOps engineer is developing an application for a company. The application needs to persist files to Amazon S3. The application needs to upload files with different security classifications that the company defines. These classifications include confidential, private, and public. Files that have a confidential classification must not be viewable by anyone other than the user who uploaded them. The application uses the IAM role of the user to call the S3 API operations. The DevOps engineer has modified the application to add a DataClassification tag with the value of confidential and an Owner tag with the uploading user's ID to each confidential object that is uploaded to Amazon S3. Which set of additional steps must the DevOps engineer take to meet the company's requirements?

A. Modify the S3 bucket's ACL to grant bucket-owner-read access to the uploading user's IAM role. Create an IAM policy that grants s3:GetObject operations on the S3 bucket when aws:ResourceTag/DataClassification equals confidential, and s3:ExistingObjectTag/Owner equals \${aws:userid}. Attach the policy to the IAM roles for users who require access to the S3 bucket.

B. Modify the S3 bucket policy to allow the s3:GetObject action when aws:ResourceTag/DataClassification equals confidential, and s3:ExistingObjectTag/Owner equals \${aws:userid}. Create an IAM policy that grants s3:GetObject operations on the S3 bucket. Attach the policy to the IAM roles for users who require access to the S3 bucket.

C. Modify the S3 bucket policy to allow the s3:GetObject action when aws:ResourceTag/DataClassification equals confidential, and aws:RequestTag/Owner equals \${aws:userid}. Create an IAM policy that grants s3:GetObject operations on the S3 bucket. Attach the policy to the IAM roles for users who require access to the S3 bucket.

D. Modify the S3 bucket's ACL to grant authenticated-read access when aws:ResourceTag/DataClassification equals confidential, and s3:ExistingObjectTag/Owner equals \${aws:userid}. Create an IAM policy that grants s3:GetObject operations on the S3 bucket. Attach the policy to the IAM roles for users who require access to the S3 bucket.

**Answer(s): B**

---

8. A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline.

A DevOps Engineer has noticed there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation, the DevOps Engineer believes the failures are due to database changes not having fully propagated before the Lambda function begins executing.

How should the DevOps Engineer overcome this?

A. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before traffic can flow to the new version of the Lambda function

B. Add an AfterAllowTraffic hook to the AppSpec file that forces traffic to wait for any pending database changes before allowing the new version of the Lambda function to respond

C. Add a BeforeInstall hook to the AppSpec file that tests and waits for any necessary database changes before deploying the new version of the Lambda function

D. Add a ValidateService hook to the AppSpec file that inspects incoming traffic and rejects the payload if dependent services, such as the database, are not yet ready

**Answer(s):** A

---

**9.** A software company wants to automate the build process for a project where the code is stored in GitHub. When the repository is updated, source code should be compiled, tested, and pushed to Amazon S3.

Which combination of steps would address these requirements? (Choose three.)

A. Add a buildspec.yml file to the source code with build instructions.

B. Configure a GitHub webhook to trigger a build every time a code change is pushed to the repository.

C. Create an AWS CodeBuild project with GitHub as the source repository.

D. Create an AWS CodeDeploy application with the Amazon EC2/On-Premises compute platform.

E. Create an AWS OpsWorks deployment with the install dependencies command.

F. Provision an Amazon EC2 instance to perform the build.

**Answer(s):** A B C

---

**10.** An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.

When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.

How should the company meet these requirements with the LEAST amount of application changes?

A. Use Amazon Redshift for the product catalog and Amazon DynamoDB tables for the customer information and purchases.

B. Use Amazon DynamoDB global tables for the product catalog and regional tables for the customer information and purchases.

C. Use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases.

D. Use Aurora for the product catalog and Amazon DynamoDB global tables for the customer information and purchases.

**Answer(s): C**

---

**11.** 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.

**Answer(s): D**

---

**12.** A DevOps Engineer needs to back up sensitive Amazon S3 objects that are stored within an S3 bucket with a private bucket policy using the S3 cross-region replication functionality. The objects need to be copied to a target bucket in a different AWS Region and account. Which actions should be performed to enable this replication? (Choose three.)

- A. Create a replication IAM role in the source account.
- B. Create a replication IAM role in the target account.
- C. Add statements to the source bucket policy allowing the replication IAM role to replicate objects.
- D. Add statements to the target bucket policy allowing the replication IAM role to replicate objects.
- E. Create a replication rule in the source bucket to enable the replication.
- F. Create a replication rule in the target bucket to enable the replication.

**Answer(s):** A D E

---

**13.** A company is using Amazon EC2 for various workloads. Company policy requires that instances be managed centrally to standardize configurations. These configurations include standard logging, metrics, security assessments, and weekly patching.

How can the company meet these requirements? (Choose three.)

- A. Use AWS Config to ensure all EC2 instances are managed by Amazon Inspector.
- B. Use AWS Config to ensure all EC2 instances are managed by AWS Systems Manager.
- C. Use AWS Systems Manager to install and manage Amazon Inspector, Systems Manager Patch Manager, and the Amazon CloudWatch agent on all instances.
- D. Use Amazon Inspector to install and manage AWS Systems Manager, Systems Manager Patch Manager, and the Amazon CloudWatch agent on all instances.
- E. Use AWS Systems Manager maintenance windows with Systems Manager Run Command to schedule Systems Manager Patch Manager tasks. Use the Amazon CloudWatch agent to schedule Amazon Inspector assessment runs.
- F. Use AWS Systems Manager maintenance windows with Systems Manager Run Command to schedule Systems Manager Patch Manager tasks. Use Amazon CloudWatch Events to schedule Amazon Inspector assessment runs.



**Answer(s):** B C F

---

**14.** A business has an application that consists of five independent AWS Lambda functions. The DevOps Engineer has built a CI/CD pipeline using AWS CodePipeline and AWS CodeBuild that builds, tests, packages, and deploys each Lambda function in sequence. The pipeline uses an Amazon CloudWatch Events rule to ensure the pipeline execution starts as quickly as possible after a change is made to the application source code.

After working with the pipeline for a few months, the DevOps Engineer has noticed the pipeline takes too long to complete. What should the DevOps Engineer implement to BEST improve the speed of the pipeline?

A. Modify the CodeBuild projects within the pipeline to use a compute type with more available network throughput.

B. Create a custom CodeBuild execution environment that includes a symmetric multiprocessing configuration to run the builds in parallel.

C. Modify the CodePipeline configuration to execute actions for each Lambda function in parallel by specifying the same runOrder.

D. Modify each CodeBuild project to run within a VPC and use dedicated instances to increase throughput.

**Answer(s):** C

---

**15.** A company is creating a software solution that executes a specific parallel-processing mechanism. The software can scale to tens of servers in some special scenarios. This solution uses a proprietary library that is license-based, requiring that each individual server have a single, dedicated license installed. The company has 200 licenses and is planning to run 200 server nodes concurrently at most.

The company has requested the following features:

A mechanism to automate the use of the licenses at scale.

Creation of a dashboard to use in the future to verify which licenses are available at any moment.

What is the MOST effective way to accomplish these requirements?

A. Upload the licenses to a private Amazon S3 bucket. Create an AWS CloudFormation template with a Mappings section for the licenses. In the template, create an Auto Scaling group to launch the servers. In the user data script, acquire an available license from the Mappings section. Create an Auto Scaling lifecycle hook, then use it to update the mapping after the instance is terminated.

B. Upload the licenses to an Amazon DynamoDB table. Create an AWS CloudFormation template that uses an Auto Scaling group to launch the servers. In the user data script, acquire an available license from the DynamoDB table. Create an Auto Scaling lifecycle hook, then use it to update the mapping after the instance is terminated.

C. Upload the licenses to a private Amazon S3 bucket. Populate an Amazon SQS queue with the list of licenses stored in S3. Create an AWS CloudFormation template that uses an Auto Scaling group to launch the servers. In the user data script acquire an available license from SQS. Create an Auto Scaling lifecycle hook, then use it to put the license back in SQS after the instance is terminated.

D. Upload the licenses to an Amazon DynamoDB table. Create an AWS CLI script to launch the servers by using the parameter `--count`, with `min:max` instances to launch. In the user data script, acquire an available license from the DynamoDB table. Monitor each instance and, in case of failure, replace the instance, then manually update the DynamoDB table.

**Answer(s): B**

---

**16.** A DevOps Engineer administers an application that manages video files for a video production company. The application runs on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. Data is stored in an Amazon RDS PostgreSQL

Multi-AZ DB instance, and the video files are stored in an Amazon S3 bucket. On a typical day, 50 GB of new video are added to the S3 bucket. The Engineer must implement a multi-region disaster recovery plan with the least data loss and the lowest recovery times. The current application infrastructure is already described using AWS CloudFormation.

Which deployment option should the Engineer choose to meet the uptime and recovery objectives for the system?

A. Launch the application from the CloudFormation template in the second region, which sets the capacity of the Auto Scaling group to 1. Create an Amazon RDS read replica in the second region. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucket. To fail over, promote the read replica as master. Update the CloudFormation stack and increase the capacity of the Auto Scaling group.

B. Launch the application from the CloudFormation template in the second region, which sets the capacity of the Auto Scaling group to 1. Create a scheduled task to take daily Amazon RDS cross-region snapshots to the second region. In the second region, enable cross-region replication between the original S3 bucket and Amazon Glacier. In a disaster, launch a new application stack in the second region and restore the database from the most recent snapshot.

C. Launch the application from the CloudFormation template in the second region, which sets the capacity of the Auto Scaling group to 1. Use Amazon CloudWatch Events to schedule a nightly task to

take a snapshot of the database, copy the snapshot to the second region, and replace the DB instance in the second region from the snapshot. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucket. To fail over, increase the capacity of the Auto Scaling group.

D. Use Amazon CloudWatch Events to schedule a nightly task to take a snapshot of the database and copy the snapshot to the second region. Create an AWS Lambda function that copies each object to a new S3 bucket in the second region in response to S3 event notifications. In the second region, launch the application from the CloudFormation template and restore the database from the most recent snapshot.

**Answer(s): A**

---

**17.** A company is using AWS CodePipeline to automate its release pipeline. AWS CodeDeploy is being used in the pipeline to deploy an application to Amazon ECS using the blue/green deployment model. The company wants to implement scripts to test the green version of the application before shifting traffic. These scripts will complete in 5 minutes or less. If errors are discovered during these tests, the application must be rolled back.

Which strategy will meet these requirements?

A. Add a stage to the CodePipeline pipeline between the source and deploy stages. Use AWS CodeBuild to create an execution environment and build commands in the buildspec file to invoke test scripts. If errors are found, use the `aws deploy stop-deployment` command to stop the deployment.

B. Add a stage to the CodePipeline pipeline between the source and deploy stages. Use this stage to execute an AWS Lambda function that will run the test scripts. If errors are found, use the `aws deploy stop-deployment` command to stop the deployment.

C. Add a hooks section to the CodeDeploy AppSpec file. Use the `AfterAllowTestTraffic` lifecycle event to invoke an AWS Lambda function to run the test scripts. If errors are found, exit the Lambda function with an error to trigger rollback.

D. Add a hooks section to the CodeDeploy AppSpec file. Use the `AfterAllowTraffic` lifecycle event to invoke the test scripts. If errors are found, use the `aws deploy stop-deployment` CLI command to stop the deployment.

**Answer(s): C**

---

**18.** A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery. Which combination of deployment strategies will meet these requirements? (Choose two.)

- A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store. Use Aurora's automatic recovery capabilities in the event of a disaster.
- B. Create an Amazon Aurora global database in two Regions as the data store. In the event of a failure, promote the secondary Region as the master for the application.
- C. Create an Amazon Aurora multi-master cluster across multiple Regions as the data store. Use a Network Load Balancer to balance the database traffic in different Regions.
- D. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Regions. Use health checks to determine the availability in a given Region. Use Auto Scaling groups in each Region to adjust capacity based on demand.
- E. Set up the application in two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on demand. In the event of a disaster, adjust the Auto Scaling group's desired instance count to increase baseline capacity in the failover Region.

**Answer(s):** B D

---

**19.** An Amazon EC2 instance is running in a Virtual Private Cloud (VPC) and needs to download an object from a restricted Amazon S3 bucket. When the DevOps engineer tries to download, the object an AccessDenied error is received.

What are the possible causes for this error? (Choose two.)

- A. The S3 bucket default encryption is enabled
- B. There is an error in the S3 bucket policy
- C. The object has been moved to Amazon Glacier
- D. There is an error in the IAM role configuration
- E. S3 versioning is enabled

**Answer(s):** B D

---

**20.** A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency.

Which actions should be taken to accomplish this? (Choose two.)

- A. Install the CloudWatch agent server side and configure the agent to upload relevant logs to CloudWatch.
- B. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and upload those segments to X-Ray during each request.
- C. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and use the X-Ray daemon to upload segments to X-Ray.
- D. Modify the on-premises application to send log information back to API Gateway with each request.
- E. Modify the on-premises application to calculate and upload statistical data relevant to the API service requests to CloudWatch metrics.

**Answer(s):** A C

---