

Configuration Manual

MSc Research Project
Cloud Computing

Paris Moore
Student ID: x14485758

School of Computing
National College of Ireland

Supervisor: Horacio Gonzalez-Velez

National College of Ireland
Project Submission Sheet
School of Computing



Student Name:	Paris Moore
Student ID:	x14485758
Programme:	Cloud Computing
Year:	2021-22
Module:	MSc Research Project
Supervisor:	Horacio Gonzalez-Velez
Submission Due Date:	15th Aug 2022
Project Title:	'Continuous Benchmarking' in DevOps to support Quality of Deployments using Amazon Web Services
Word Count:	2042
Page Count:	18

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	Paris Moore
Date:	15th September 2022

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	<input type="checkbox"/>
Attach a Moodle submission receipt of the online project submission , to each project (including multiple copies).	<input type="checkbox"/>
You must ensure that you retain a HARD COPY of the project , both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	<input type="checkbox"/>

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

'Continuous Benchmarking' in DevOps to support Quality of Deployments using Amazon Web Services

Paris Moore
x14485758

1 Introduction

This configuration manual will help its readers to understand the system requirements, setup, software and install specifications that were used in this research. Also, this manual includes detailed explanation of the steps needed to follow when implementing this research project. The proof-of-concept pipeline is designed and built using Amazon Web Services' (AWS), which automatically deploys the system release, runs one or more benchmarks, collects and analyzes results, and decides whether the release fulfils pre-defined Quality of Deployment (QoD) goals.

Prerequisites

- AWS Cloud Knowledge
- Basic Command Line Knowledge
- AWS IAM access to create roles and add policies to users
- Have Python installed to be able to run the application locally.
- Have Git installed locally.

2 Before you begin

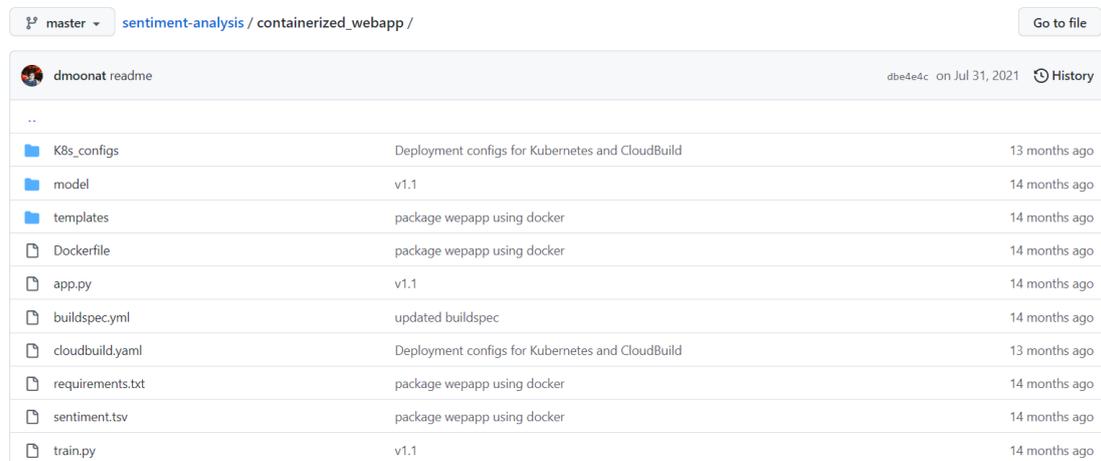
2.1 Software Installation

- Install AWS CLI locally following the steps for your Operating System
<https://docs.aws.amazon.com/cli/v1/userguide/install-windows.html>
- Install Docker locally following the steps for your Operating System
<https://docs.docker.com/get-docker/>

2.2 The Application

An application is required for analysis purposes and to test deployments on the pipeline. An open source Twitter Sentiment Analysis application is chosen. The app is built on the python flask framework with machine learning models developed to perform the predictions. The application will be containerized using docker and deployed using AWS Elastic

Container Service (ECS). Inside directory 'sentiment-analysis/containerized_webapp/' in the Application source file is a .Dockerfile and buildspec.yml files which will be very important when deploying the application on AWS Cloud.



File/Folder	Description	Last Modified
..		
KSs_configs	Deployment configs for Kubernetes and CloudBuild	13 months ago
model	v1.1	14 months ago
templates	package wepapp using docker	14 months ago
Dockerfile	package wepapp using docker	14 months ago
app.py	v1.1	14 months ago
buildspec.yml	updated buildspec	14 months ago
cloudbuild.yml	Deployment configs for Kubernetes and CloudBuild	13 months ago
requirements.txt	package wepapp using docker	14 months ago
sentiment.tsv	package wepapp using docker	14 months ago
train.py	v1.1	14 months ago

GitHub Repository to download or clone application code:

https://github.com/dmoonat/sentiment-analysis/tree/master/containerized_webapp

Once the application has been downloaded and saved. The application can be ran locally by running the following commands in the 'sentiment-analysis/containerized_webapp/' directory.

- pip install -r requirements.txt
- python app.py

3 Deploying the Application on AWS

3.1 Docker

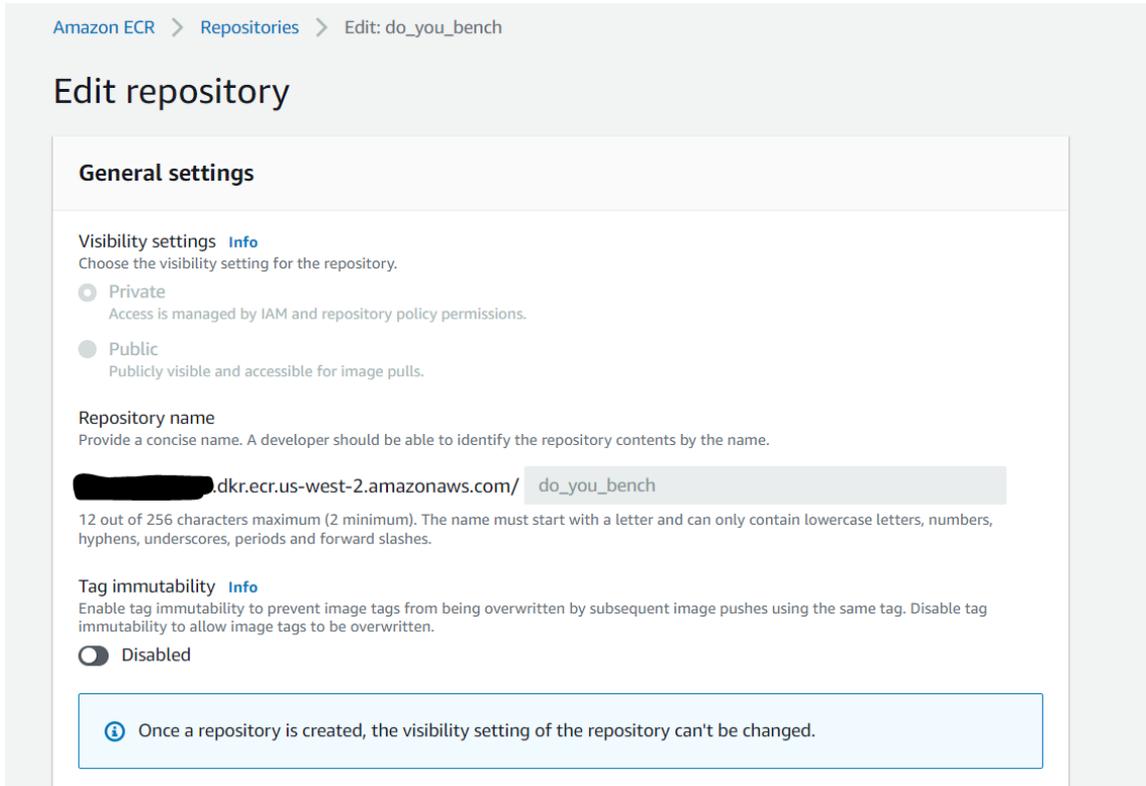
Once you have installed Docker, you can verify the installation is successful by running the 'docker' command on the terminal/command prompt. Its time to dockerize the Sentiment Analysis Application. The dockerfile already includes the necessary commands for docker to build an image.

- docker build -t final_project_2022:v1 -f Dockerfile .
- Check an image has been built using the 'docker images' command
- Test the image by running it using this command:
docker run -p 5000:5000 -t -i final_project_2022:v1
- Go to the IP of the docker displayed to check that the app is running on a container.

3.2 Elastic Container Registry

AWS ECR is a registry to store and manage container images, similar to DockerHub.

- Login to AWS management console → ECR
- Create a repository for the application



- To communicate with Amazon services from your local machine, you must install the AWS CLI and configure your profile using the terminal. For NCI students, this must be done via SSO. For those with admin access, this can be achieved by generating secret access keys within your AWS account.

```

:~$ aws configure
AWS Access Key ID [None]: 
AWS Secret Access Key [None]: 
Default region name [None]: us-east-1
Default output format [None]:

```

- Once AWS is configured, next you need to authenticate to your ECR repo using this command with your 'region' 'profile' and 'aws account ID' defined:

```

aws ecr get-login-password --region 'region' --profile 'profile'
| docker login --username AWS --password-stdin 'aws_account_id'.dkr.ecr.'region'.amazonaws.com

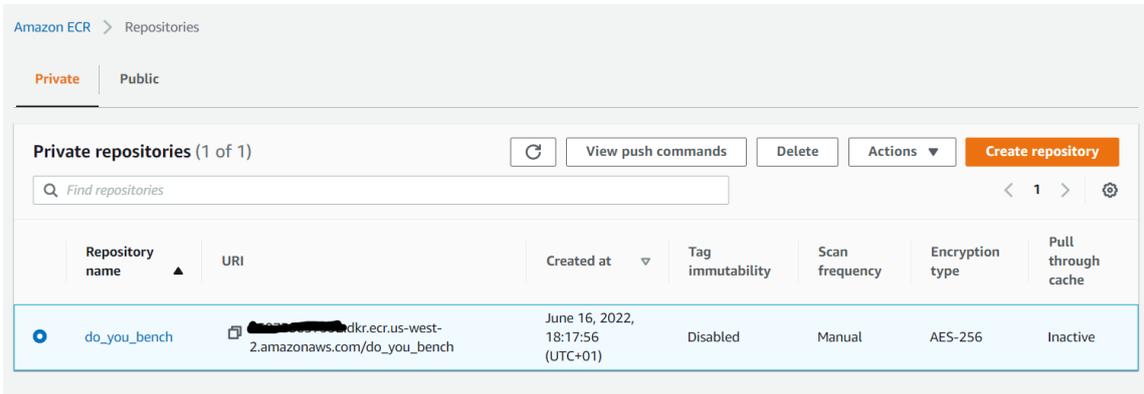
```

Further supporting documentation can be found here:

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/getting-started-cli.html>

- Once Authenticated, follow the ECR push commands from your local terminal to upload a docker image to the repo.

- You can confirm that your image has been pushed by checking the ECR console.



```
C:\Users\rbannon\OneDrive - Irish Mapping & GIS Solutions Ltd\Desktop\Paris College\Sentiment Analysis\sentiment-analysis\containerized_webapp> aws ecr get-login-password --region us-west-2 --profile s3485738-P-2021-250738637992 | docker login --username AWS --password-stdin 250738637992.dkr.ecr.us-west-2.amazonaws.com
WARNING! Your password will be stored unencrypted in C:\Users\rbannon\.docker\config.json.
Configure a credential helper to remove this warning. See
https://docs.docker.com/engine/reference/commandline/login/#credentials-store

Login Succeeded

C:\Users\rbannon\OneDrive - Irish Mapping & GIS Solutions Ltd\Desktop\Paris College\Sentiment Analysis\sentiment-analysis\containerized_webapp> docker build -t do_you_bench .
[+] Building 23.7s (13/13) [DN]298D
=> [internal] load build definition from Dockerfile
=> => transferring Dockerfile: 32B
=> [internal] load .dockerignore
=> => transferring context: 48
=> resolve image config for docker.io/docker/dockerfile:1
=> CACHED docker image //docker.io/docker/dockerfile:1@sha256:443a8dc42118369bc5062bc0800c944f40bd010136a81f51176d36b6c2
=> [internal] load .dockerignore
=> [internal] load build definition from Dockerfile
=> [internal] load metadata for docker.io/library/python:3.8-slim-buster@sha256:75c0bc0f6bc6223c3c0941f8623489bc3a6262f91f48f18270945c0d7003
=> [1/4] FROM docker.io/library/python:3.8-slim-buster@sha256:75c0bc0f6bc6223c3c0941f8623489bc3a6262f91f48f18270945c0d7003
=> [internal] load build context
=> => transferring context: 32B
=> CACHED [2/4] WORKDIR /app
=> CACHED [3/4] COPY . /app
=> CACHED [4/4] RUN pip install -r requirements.txt
=> exporting to image
=> => writing image sha256:858051a6e705e48208cc1f96277951f36e1301f88148e775261a1914841e9
=> => pushing to docker.io/library/do_you_bench

Use 'docker scan' to run Snyk tests against images to find vulnerabilities and learn how to fix them

C:\Users\rbannon\OneDrive - Irish Mapping & GIS Solutions Ltd\Desktop\Paris College\Sentiment Analysis\sentiment-analysis\containerized_webapp> docker tag do_you_bench:latest 250738637992.dkr.ecr.us-west-2.amazonaws.com/do_you_bench:latest
The push refers to repository [250738637992.dkr.ecr.us-west-2.amazonaws.com/do_you_bench]
76c1f3642fab: Pushed
a8fa7187c76a: Pushed
686ad19e2bed: Pushed
9f8ed4ffacbc: Pushed
f99d78f4d2d7: Pushed
42fac8006c35: Pushed
86a9e23f9e2: Pushed
18ebc646ae2: Pushed
latest digest: sha256:4f6bd34249b5b6c0903954b564cfab5005e2a6a3e18182391edf545ec0b74 size: 2000
```

3.3 Elastic Container Service

Amazon ECS is a container orchestration platform developed by Amazon, that helps schedule and orchestrate containers across a group of servers. The two major ECS components are Tasks and Services.

- First step is to create a cluster using ECS. Select 'Networking Only' and then create.

Amazon ECS

- Clusters**
- Task Definitions
- Account Settings

Amazon EKS

- Clusters

Amazon ECR

- Repositories

AWS Marketplace

- Discover software
- Subscriptions [↗](#)

Cluster : sentiment-analysis-cluster

Get a detailed view of the resources on your cluster.

Cluster ARN arn:aws:ecs:us-west-2:250738637992:cluster/sentiment-analysis-cluster

Status ACTIVE

Registered container instances 0

Pending tasks count 0 Fargate, 0 EC2, 0 External

Running tasks count 3 Fargate, 0 EC2, 0 External

Active service count 1 Fargate, 0 EC2, 0 External

Draining service count 0 Fargate, 0 EC2, 0 External

Services | **Tasks** | ECS Instances | Metrics | Scheduled Tasks | Tags | Capacity Providers

Create | Update | Delete | Actions

- Next you must create a 'Task Definition' and choose 'Fargate' as launch type. Give the task definition a name, specify task Memory and CPU needed to run the task.

Task size

The task size allows you to specify a fixed size for your task. Task size is required for tasks using the Fargate launch type and is optional for the EC2 or External launch type. Container level memory settings are optional when task size is set. Task size is not supported for Windows containers.

Task memory (GB) 1GB
The valid memory range for 0.25 vCPU is: 0.5GB - 2GB.

Task CPU (vCPU) 0.25 vCPU
The valid CPU range for 1GB memory is: 0.25 vCPU - 0.5 vCPU.

Task memory maximum allocation for container memory reservation
0 1024 shared of 1024 MiB

Task CPU maximum allocation for containers
0 256 shared of 256 CPU units

Container Definitions

Add container

Container ...	Image	Hard/Soft ...	CPU Units	GPU	Inference A...	Essential
No results						

- On the same page, click 'Add container' and provide a name and the image URI (available from your ECR repo). As per how the application is developed, specify the port mappings to 5000 so containers can send or receive traffic.

Container definitions

Container Name	Image	CPU Units	GPU	Inference Ac...	Hard/Soft memory limits (MiB)	Essential
sentiment-analysis-container	25073863799...	0			--/--	true

Port Mappings			Mount Points		
Host Port	Container Port	Protocol	Container Path	Source Volume	Read only
5000	5000	tcp	No mount points		
			Volumes from		
			Source Container	Read only	

- Click on Actions and Run Task, select launch type as Fargate. Select the VPC and subnet from the dropdown and click Run Task.
- Our Task is created successfully, now we need to add an inbound rule to the security group to access our application on port 5000.
- Click on the created task → click 'ENI Id' → network interface page. Click on the network checkbox and scroll down to 'Security groups'.

Network interfaces (1/1) info

Search

Network interface ID = eni-04f1125caf9850890

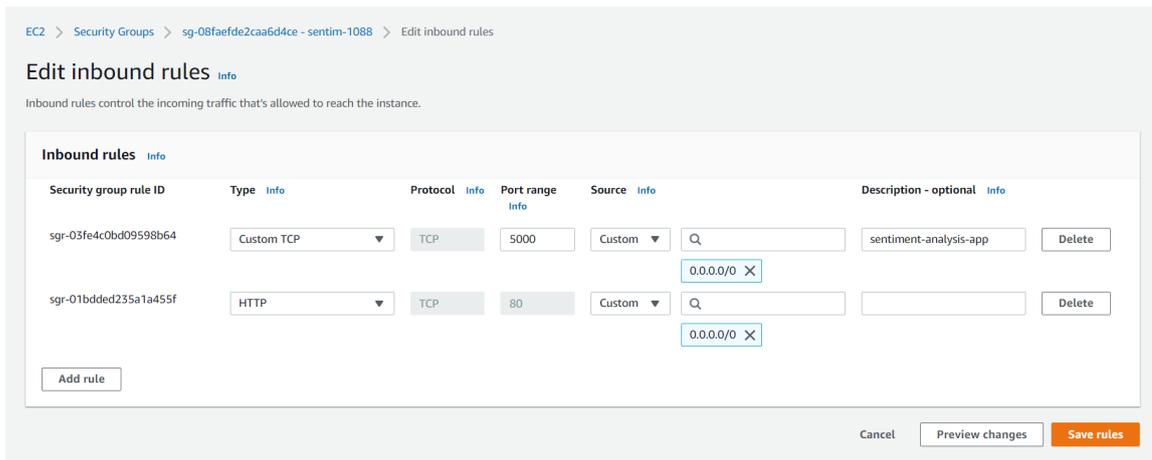
Name	Network interface ID	Subnet ID	VPC ID	Availability Zone	Security group n...
-	eni-04f1125caf9850890	subnet-00204490772be421	vpc-0b3bde132ef075470	us-west-2b	sentim-1088

Network interface: eni-04f1125caf9850890

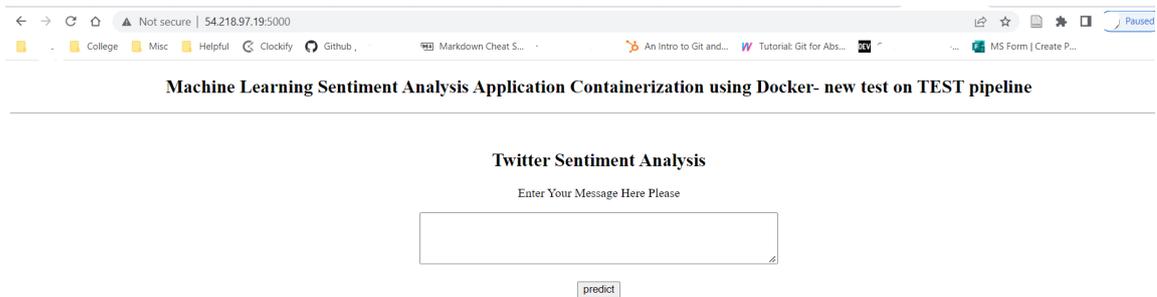
Network interface details

Network interface ID eni-04f1125caf9850890	Name -	Description arn:aws:ecs:us-west-2:250738637992:attachment/21036467-cd19-4f5b-a49e-bc4c18580ed6
Network interface status In-use	Interface type Elastic network interface	Security groups sg-08faefde2caa6d4ce (sentim-1088)
VPC ID vpc-0b3bde132ef075470	Subnet ID subnet-00204490772be421	Availability Zone us-west-2b

- Now click on Edit inbound rules and click Add rule. Add a custom TCP rule with port 5000 and source to be 0.0.0.0/0 (to make the application accessible through all IPs) and click Save rules.



- Go back to the task page and using the public IP and port 5000, you can access the sentiment analysis application from the browser.



3.4 EC2 Load Balancer

Next, we create an Application load balancer (ALB) for our application. A load balancer serves as the single point of contact for clients and distributes incoming application traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. This increases the availability and scalability of the application.

- Go to EC2 and click on Load Balancers. Choose Application load Balancer and Create.
- Add a name, VPC and availability zones.
- Create a new security group, add a name to Target group, for Target type select IP and then click Next.

VPC	vpc-0b3bde132ef075470 ↗
Availability Zones	subnet-002044907772be421 - us-west-2b ↗ IPv4 address: Assigned by AWS subnet-0dd58493a4229a46f - us-west-2a ↗ IPv4 address: Assigned by AWS Edit subnets
Hosted zone	Z1H1FL5HABSF5
Creation time	June 21, 2022 at 6:54:23 PM UTC+1

Security

Security groups	sg-08faefde2caa6d4ce, sentim-1088 • 2022-06-22T15:51:20.663Z
------------------------	--

[Edit security groups](#)

- Once created, note the DNS name, which is the public address where the application is accessed from the browser.

3.5 AWS Fargate

We will use the same task definition we already created in ECS to create a Fargate Service. AWS Fargate is a serverless compute engine for Amazon ECS that runs containers without requiring us to worry about the underlying infrastructure.

- Go to Task Definitions in Amazon ECS → tick the radio button corresponding to the existing task definition → click Actions → Create Service → Choose Fargate.
- Leave all other options as is, including rolling update as deployment type.
- Choose the same subnets as the one configured in the load balancer.
- Choose ALB as the load balancer type and add the already created ALB to it.
- Select the target group created in the ALB and click Create Service.
- Go back to the EC2 ALB and click on security groups; add a 'Custom TCP' rule with port '5000' under inbound rules. This is the internal port that the flask app is set to run on.
- Visiting the ALB DNS url previously noted, you can check the application is running on the fargate service.



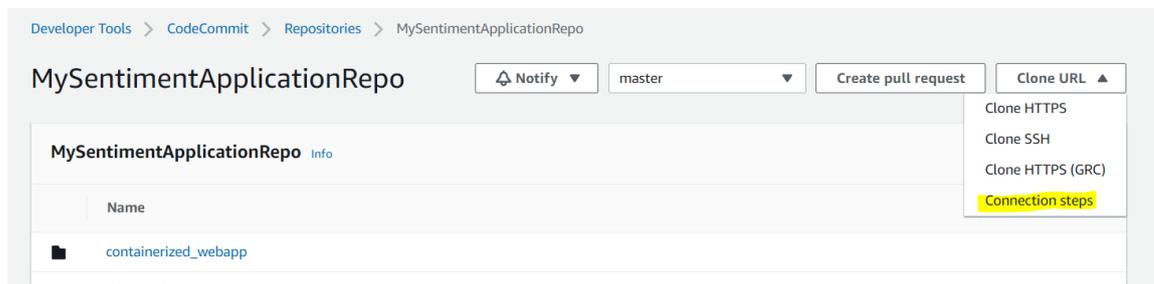
DNS URL to Running application: sentiment-analysis-app-balancer-1898041068.us-west-2.elb.amazonaws.com

4 Configuring the Pipeline

4.1 CodeCommit

AWS CodeCommit is used to store the applications artifacts. It is Amazon's version of Github.

- Click Create Repository and fill in the required details.
- Once created, click on the repo and from the drop-down choose 'Connection Steps'

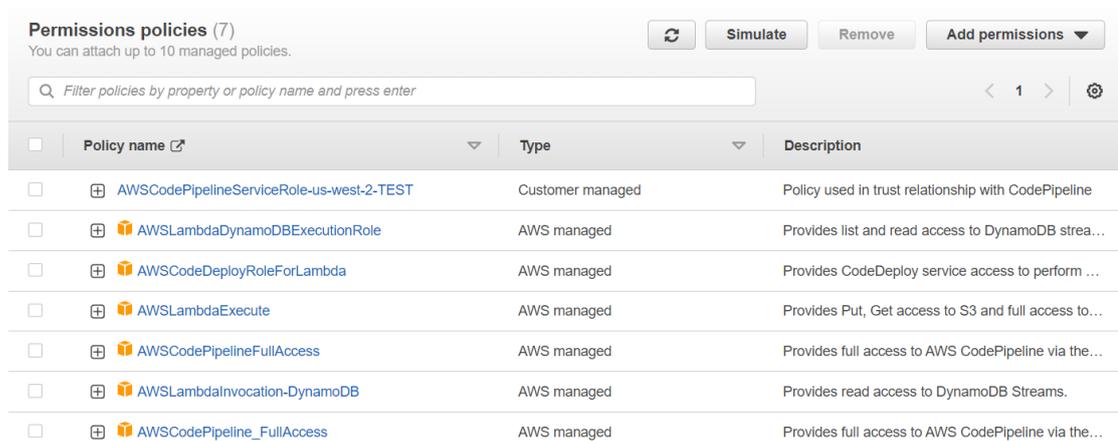


- You must still have an active session on the AWS CLI for the next steps. Otherwise, refresh the session by re-running 'aws configure'
- Next, follow the steps as outlined on the CodeCommit console to push the local application folder to the CodeCommit repository.

4.2 CodePipeline

AWS CodePipeline is used to configure and build the CI/CD/CB pipeline.

- Go to CodePipeline using the AWS Console. Click Create Pipeline.
- Provide a name and take note of the pipeline role name created for you. Click next.
- Go to IAM and add the following policies to the CodePipeline role:

The screenshot shows the 'Permissions policies (7)' section in the AWS IAM console. It includes a search bar with the placeholder text 'Filter policies by property or policy name and press enter'. Below the search bar is a table with the following columns: 'Policy name', 'Type', and 'Description'. The table lists seven policies:

Policy name	Type	Description
AWSCodePipelineServiceRole-us-west-2-TEST	Customer managed	Policy used in trust relationship with CodePipeline
AWSLambdaDynamoDBExecutionRole	AWS managed	Provides list and read access to DynamoDB stre...
AWSCodeDeployRoleForLambda	AWS managed	Provides CodeDeploy service access to perform ...
AWSLambdaExecute	AWS managed	Provides Put, Get access to S3 and full access to...
AWSCodePipelineFullAccess	AWS managed	Provides full access to AWS CodePipeline via the...
AWSLambdaInvocation-DynamoDB	AWS managed	Provides read access to DynamoDB Streams.
AWSCodePipeline_FullAccess	AWS managed	Provides full access to AWS CodePipeline via the...

4.2.1 Source

- Next we configure the source stage. Choose AWS CodeCommit as the action provider.
- Choose your application repository in CodeCommit.
- Choose 'master' as branch name. This is what will become your release trigger.
- Keep CloudWatch logs enabled.

The screenshot shows the 'Source' configuration page in the AWS CodePipeline console. It includes the following sections:

- Action name:** A text input field containing 'Source'.
- Action provider:** A dropdown menu set to 'AWS CodeCommit'.
- Repository name:** A search input field containing 'MySentimentApplicationRepo'.
- Branch name:** A search input field containing 'master'.
- Change detection options - optional:** Two radio button options: 'Amazon CloudWatch Events (recommended)' (selected) and 'AWS CodePipeline'.
- Output artifact format - optional:** Two radio button options: 'CodePipeline default' (selected) and 'Full clone'.

4.2.2 Build

- The next stage is build. Choose CodeBuild as Build Provider and click 'Create Project'.

The screenshot shows the 'Build - optional' configuration page in the AWS CodePipeline console. It includes the following sections:

- Build provider:** A dropdown menu set to 'AWS CodeBuild'.
- Region:** A dropdown menu set to 'US West (Oregon)'.
- Project name:** A search input field followed by a 'Create project' button with an external link icon.
- Environment variables - optional:** A section with a link to 'Learn more' and an 'Add environment variable' button.
- Build type:** Two radio button options: 'Single build' (selected) and 'Batch build'.

At the bottom of the page, there are four buttons: 'Cancel', 'Previous', 'Skip build stage', and 'Next'.

- Configure the build project by adding a name, choose 'custom image', environment type as 'Linux' and specify the Amazon ECR repo from the dropdown.
- Ensure you check the box that gives privileged access for CodeBuild to build docker images on your behalf.
- Add the following environment variables under Additional configuration:

Environment variables

Name	Value	Type	
AWS_ACCOUNT_ID	250736637532	Plaintext	Remove
IMAGE_REPO_NAME	do_you_bench	Plaintext	Remove
IMAGE_TAG	latest	Plaintext	Remove

[Add environment variable](#)

- Under 'Logs', ensure both cloudwatch and S3 are checked. You will need to specify what bucket in S3 you wish to use.
- Leave everything else as default and click create.
- Go to IAM and add the following policies to the CodeBuild role that has just been created:

The screenshot shows the AWS IAM console 'Permissions policies' page. It lists seven policies attached to the role:

Policy name	Type	Description
CodeBuildBasePolicy-sentiment_analysis_x14485758-...	Customer managed	Policy used in trust relationship with CodeBuild
CodeBuildBasePolicy-TEST-us-west-2	Customer managed	Policy used in trust relationship with CodeBuild
CodeBuildCloudWatchLogsPolicy-sentiment_analysis_x...	Customer managed	Policy used in trust relationship with CodeBuild
CodeBuildCloudWatchLogsPolicy-TEST-us-west-2	Customer managed	Policy used in trust relationship with CodeBuild
CodeBuildVpcPolicy-sentiment_analysis_x14485758-us...	Customer managed	Policy used in trust relationship with CodeBuild
AmazonEC2ContainerRegistryPowerUser	AWS managed	Provides full access to Amazon EC2 Container R...
EC2InstanceProfileForImageBuilderECRContainerB...	AWS managed	EC2 Instance profile for building container image...

Permissions boundary - (not set)

4.2.3 Deploy

- For deploy provider, chose Amazon ECS, cluster and service name. Include the name of the image definitions file as 'images.json' which is a command in the build.yml file that creates this file during the build process.

```

#Required sequence. Represents the commands CodeBuild runs during each phase of the build.
phases:
  pre_build:
    commands:
      - echo Logging in to Amazon ECR...
      - CODEBUILD_RESOLVED_SOURCE_VERSION="${CODEBUILD_RESOLVED_SOURCE_VERSION:-$IMAGE_TAG}"
      - IMAGE_TAG=$(echo $CODEBUILD_RESOLVED_SOURCE_VERSION | cut -c 1-7)
      - echo image_tag $IMAGE_TAG
      - REPO="$AWS_ACCOUNT_ID.dkr.ecr.$AWS_DEFAULT_REGION.amazonaws.com"
      - IMAGE_URI="$AWS_ACCOUNT_ID.dkr.ecr.$AWS_DEFAULT_REGION.amazonaws.com/$IMAGE_REPO_NAME:$IMAGE_TAG"
      - echo Repository $REPO
      - docker login -u AWS -p $(aws ecr get-login-password --region $AWS_DEFAULT_REGION) $REPO
  build:
    commands:
      - echo Build started on `date`
      - echo Building the Docker image...
      - docker build -t $IMAGE_URI .
  post_build:
    commands:
      - bash -c "if [ /"$CODEBUILD_BUILD_SUCCEEDING/" == /"0/" ]; then exit 1; fi"
      - echo Build stage successfully completed on `date`
      - echo Pushing the Docker image...
      - docker push $IMAGE_URI
      - printf '[{"name":"sentiment-analysis-container","imageUri":"%s"}]' "$IMAGE_URI" > images.json
artifacts:
  files: images.json

```

- Click next and Create Pipeline.

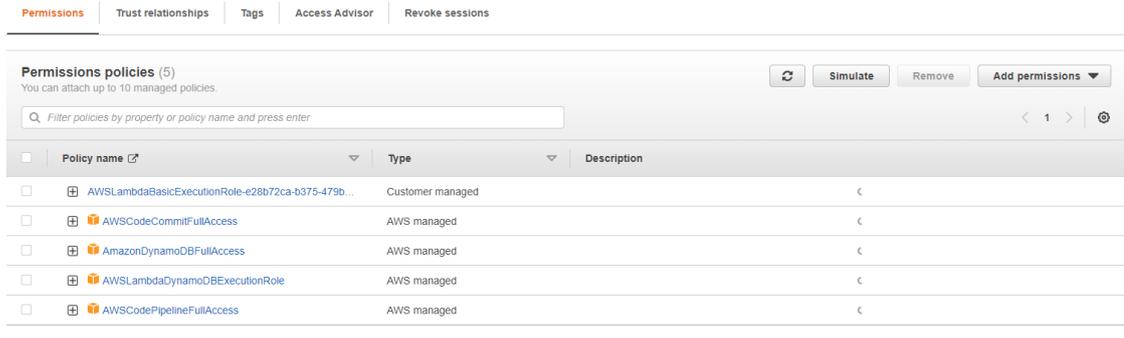
5 Implementing Benchmarking

Once the pipeline is built and the application has been deployed, its time to develop benchmarking using Lambda. AWS Lambda is an event-driven, serverless computing platform. Two additional stages will be created in the pipeline, both which will trigger lambda functions. The first benchmark will take place after the source stage and the second will be after the build stage.

Benchmarking after the Source Stage

This lambda function is dependent on event data from CodeCommit in order to perform the benchmarks.

- Go to Lambda and create a new function. Choose 'Author from scratch'. Create a new role and add the following permissions:



- The python code for this function is below:

```

1 |
2 import json
3 import boto3
4 from datetime import datetime
5
6 print('Loading function')
7
8 def lambda_handler(event, context):
9     n = 3
10    dynamodb = boto3.resource('dynamodb')
11    table = dynamodb.Table("BuildTimes")
12    pipeline = boto3.client('codepipeline')
13    codecommit = boto3.client('codecommit')
14
15    records = table.scan()
16    records = records['Items']
17    dynamoResponse = sorted(records, key=lambda x: datetime.strptime(x["Timestamp"], '%Y-%m-%d %H:%M:%S.%f'), reverse=True)[0]
18    lastDeployment = dynamoResponse['Timestamp']
19    print("Last Deployment: "+lastDeployment)
20
21    pullRequestsresponse = codecommit.list_pull_requests(
22        repositoryName='MySentimentApplicationRepo'
23    )
24
25    list = []
26    for i in pullRequestsresponse['pullRequestIds']:
27        pullRequestresponse = codecommit.get_pull_request(
28            pullRequestId=i
29        )
30        timestamp = pullRequestresponse['pullRequest']['lastActivityDate'].strftime('%Y-%m-%d %H:%M:%S.%f')
31        if(timestamp > lastDeployment):
32            if(pullRequestresponse['pullRequest']['pullRequestTargets'][0]['mergeMetadata']['isMerged']):
33                print("Merge since last deployment")
34                print(timestamp)
35                list.append(pullRequestresponse['pullRequest']['title'])
36
37    print(list)
38    if(len(list) >= n):
39        print('the release can continue')
40        response = pipeline.put_job_success_result(
41            jobId=event['CodePipeline.job']['id']
42        )
43    else:
44        print('the release criteria has not been met')
45        print('failing release')
46        response = pipeline.put_job_failure_result(
47            jobId=event['CodePipeline.job']['id'],
48            failureDetails={'message': 'Stopping Release', 'type': 'JobFailed'}
49        )
50
51    return response
52

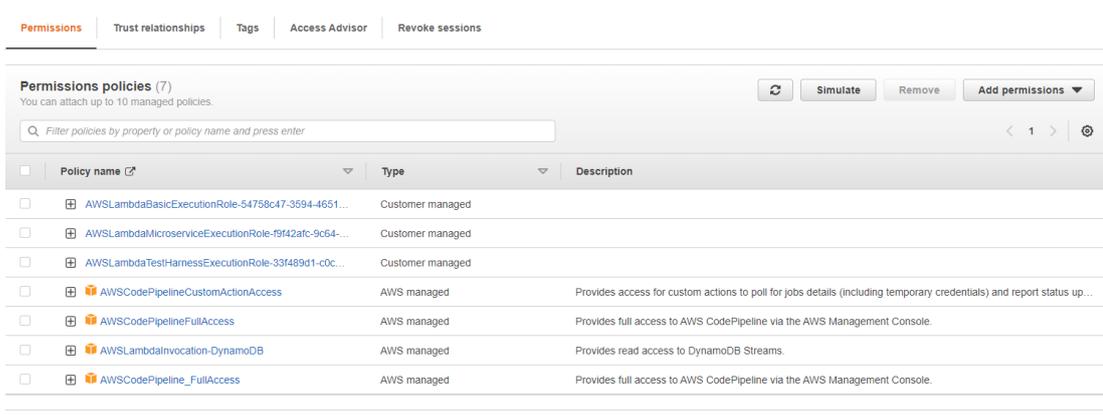
```

- Boto3 is an AWS SDK that provides a Python API for AWS infrastructure services. We use this SDK to query event data from CodeCommit. 16

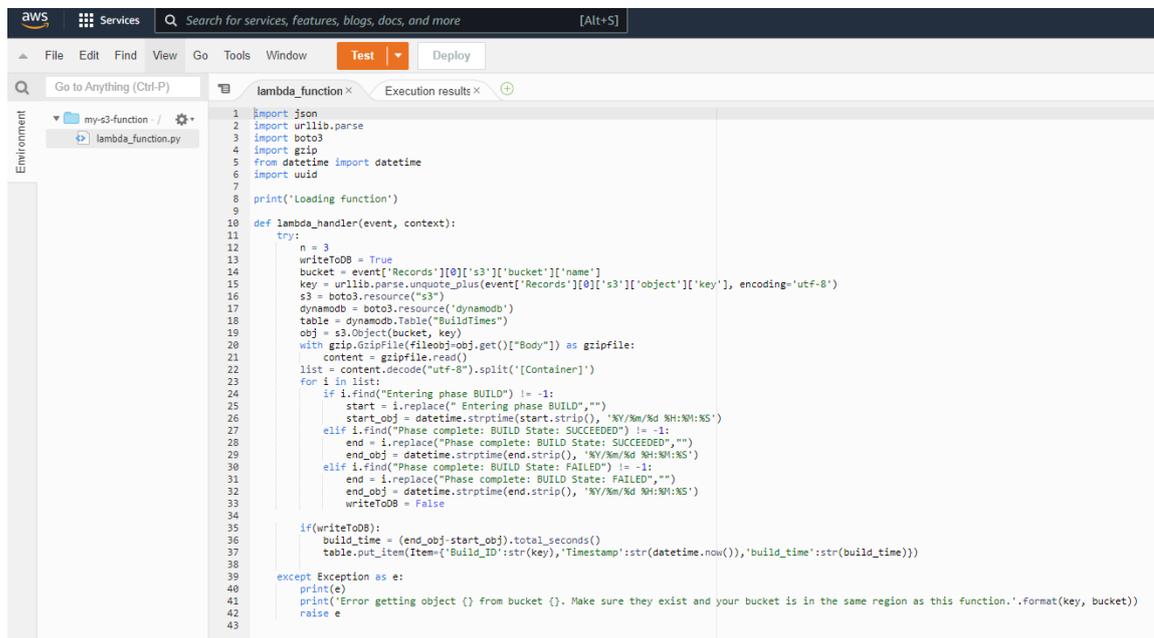
Benchmarking after the Build Stage

This lambda function is dependent on log data from CodeBuild in order to perform the benchmarks.

- The build project was setup to send the logs to an s3 bucket. By default, these appear as zipped files in s3 so in order to be able to review the logs and target certain key events, we will create a lambda function to unzip the log files from s3 and send our targeted attributes to a table in DynamoDB.
- Go to Lambda and create a new function. Choose 'Author from scratch'. Create a new role and add the following permissions:



- The python code for this function is below:



- To trigger this function, we add a lambda trigger to the S3 bucket where the build logs are being sent. Therefore, everytime a build completes, the log files are sent from cloudwatch to s3.

Event notifications (1) Edit Delete Create event notification

Send a notification when specific events occur in your bucket. [Learn more](#)

<input type="checkbox"/>	Name	Event types	Filters	Destination type	Destination
<input type="checkbox"/>	S3ObjectPut	Put	-	Lambda function	my-s3-function

Amazon EventBridge Edit

For additional capabilities, use Amazon EventBridge to build event-driven applications at scale using S3 event notifications. [Learn more](#) or [see EventBridge pricing](#)

Send notifications to Amazon EventBridge for all events in this bucket

Off

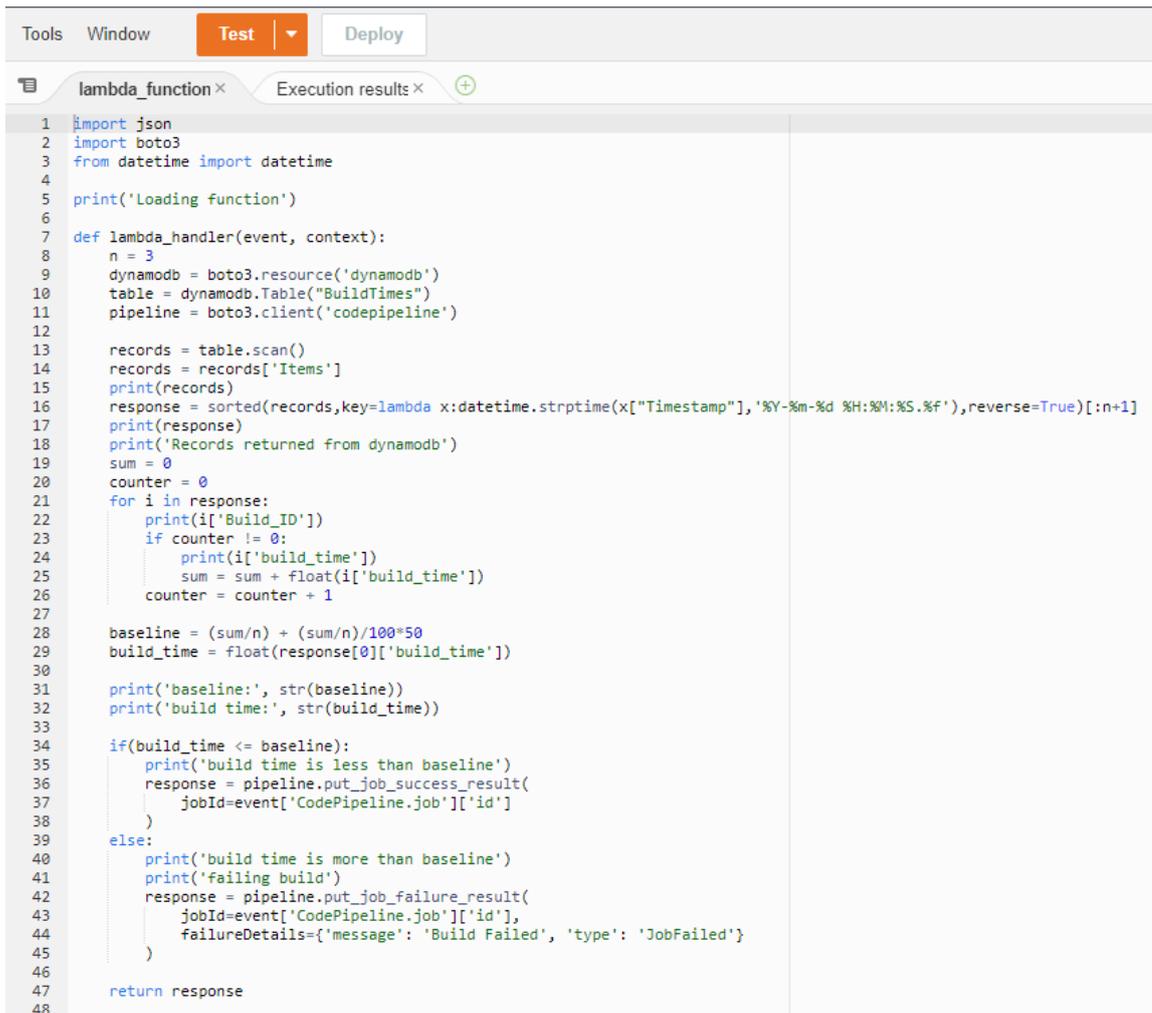
- Once they land in s3, our lambda function is triggered to get the data we need and store them in our dynamo table:

Items returned (50) Refresh Actions Create item

< 1 ... > Settings Fullscreen

<input type="checkbox"/>	Build_ID	Timestamp	build_time
<input type="checkbox"/>	c002f23a-5f7e-4d9c-a2c4-cb8d1b5558e3.gz	2022-08-09 15:34:39.184255	32.0
<input type="checkbox"/>	29cc113f-634a-417f-9f02-8287eb0419fc.gz	2022-08-09 17:05:07.042794	32.0
<input type="checkbox"/>	6e14e959-5dfc-4f51-bd9c-10f950dad6a2.gz	2022-08-08 15:40:05.529062	52
<input type="checkbox"/>	e5b0eeba-8d1f-430d-855b-27a5a2ee7937.gz	2022-08-09 10:51:06.320893	32.0
<input type="checkbox"/>	69da8f30-87ac-48dd-a62e-ddad2d1caf49.gz	2022-08-08 16:29:10.557199	53
<input type="checkbox"/>	da293d9b-76bb-48a3-b660-64852045a1a3.gz	2022-08-01 14:10:32.271023	53

- Next we develop the logic to perform the benchmarking, the python code for this is below:



```
1 import json
2 import boto3
3 from datetime import datetime
4
5 print('Loading function')
6
7 def lambda_handler(event, context):
8     n = 3
9     dynamodb = boto3.resource('dynamodb')
10    table = dynamodb.Table("BuildTimes")
11    pipeline = boto3.client('codepipeline')
12
13    records = table.scan()
14    records = records['Items']
15    print(records)
16    response = sorted(records, key=lambda x: datetime.strptime(x["Timestamp"], '%Y-%m-%d %H:%M:%S.%f'), reverse=True)[:n+1]
17    print(response)
18    print('Records returned from dynamodb')
19    sum = 0
20    counter = 0
21    for i in response:
22        print(i['Build_ID'])
23        if counter != 0:
24            print(i['build_time'])
25            sum = sum + float(i['build_time'])
26            counter = counter + 1
27
28    baseline = (sum/n) + (sum/n)/100*50
29    build_time = float(response[0]['build_time'])
30
31    print('baseline:', str(baseline))
32    print('build time:', str(build_time))
33
34    if build_time <= baseline:
35        print('build time is less than baseline')
36        response = pipeline.put_job_success_result(
37            jobId=event['CodePipeline.job']['id']
38        )
39    else:
40        print('build time is more than baseline')
41        print('failing build')
42        response = pipeline.put_job_failure_result(
43            jobId=event['CodePipeline.job']['id'],
44            failureDetails={'message': 'Build Failed', 'type': 'JobFailed'}
45        )
46
47    return response
48
```

- Boto3 is an AWS SDK that provides a Python API for AWS infrastructure services. We use this SDK to communicate with other services such as S3 and dynamo.

Add Benchmarking Stages to the Pipeline

Once the lambda functions have been developed, its time to add two new benchmarking stages to our pipeline.

- Go to CodePipeline → Click into your pipeline → Click Edit → Add Stage.
- Name your action and choose AWS Lambda as Action provider and choose the relevant function name from lambda:

Action name
Choose a name for your action

Release-Criteria-Check

No more than 100 characters

Action provider

AWS Lambda

Region

US West (Oregon)

Input artifacts
Choose an input artifact for this action. [Learn more](#)

Add

No more than 100 characters

Function name
Choose a function that you have already created in the AWS Lambda console. Or create a function in the AWS Lambda console and then return to this task.

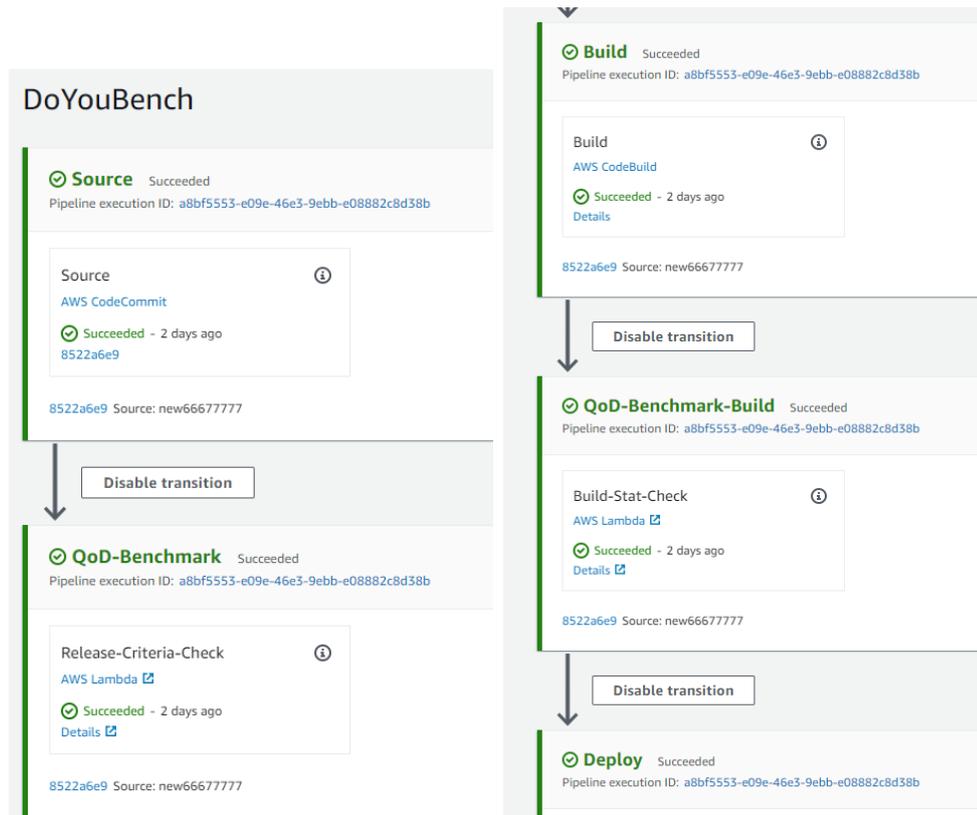
MyLambdaFunctionforCodeCommit

Function name contains only letters, numbers, hyphens, or underscores with no spaces. This does not include the function alias or function ARN.

User parameters - optional
This string will be used in the event data parameter passed to the handler in AWS Lambda.

Variable namespace - optional
Choose a namespace for the output variables from this action. You must choose a namespace if you want to use the variables this action produces in your configuration. [Learn more](#)

- Repeat the above for the second benchmarking stage. Once you have done so, your pipeline should look something like this:

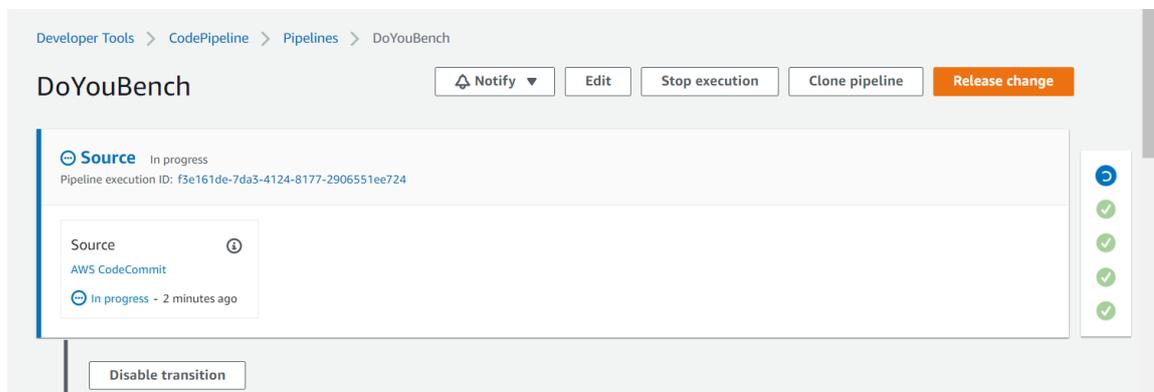


NB: The lambda functions code should be updated with suitable endpoints for your project such as CodeCommit repo, dynamo table, S3 bucket etc.

6 Running Deployments

Now its time to use the pipeline to deploy application changes.

- Make a code change to a html file and push the changes to the master branch in your CodeCommit repo.
- `git add .`
- `git commit -m " "`
- `git push`
- The pipeline release should start instantly.



- Both benchmarking functions include logic that will send a response back to the pipeline to say the benchmarks has pass or not. This will determine whether the application release can proceed to the next stage of the pipeline. If the criteria was not met, a response is sent to 'fail' that stage of the pipeline, preventing the cycle from proceeding to the next stage.
- The criteria for when each benchmark should pass and should fail is discussed in detail in the research paper.