

## Contents

Overview .....	1
Continuous Integration (CI).....	2
Continuous Delivery (CD).....	2
Purpose of the document .....	3
Scope of the document.....	3
Environment information .....	4
Architecture .....	8
Design Specification .....	8
Resources .....	12
<a href="https://www.ukessays.com/essays/computer-science/ci-and-cd-in-aws-code-pipeline.php">https://www.ukessays.com/essays/computer-science/ci-and-cd-in-aws-code-pipeline.php</a> .....	12

## Overview

This project is an introduction on how to use Amazon Web Services (AWS) in DevOps. AWS resources can decrease time to market and reduce costs for companies. The paper will discuss a specific tool called AWS code pipeline, it is a relatively new tool which was released in July 2015.

We used AWS CodePipeline to automate your software release process, allowing us to rapidly release new features to users. With CodePipeline, we can quickly iterate on feedback and get new features to the users faster.

Automating the build, test, and release process allowed us to test each code change and catch bugs while they are small and simple to fix quickly and easily. We can ensure the quality of our application or infrastructure code by running each change through our staging and release process.

In today's world applications must evolve quickly for customers. Improving and releasing software at a fast pace to customer needs to be at the core of every business. Making time and agility to market essential to maintaining your competitive advantage. Companies that can rapidly deliver updates to applications and services, innovate and change faster to adapting changing markets which give better results to business and customers.

With AWS code pipeline, companies can deliver value to their customers quickly and safely.

## Continuous Integration (CI)

Continuous integration (CI) is a software development practice where developers regularly merge their code changes into a central repository, after which automated builds and tests are run. CI most often refers to the build or integration stage of the software release process and requires both an automation component (for example a CI or build service) and a cultural component (for example learning to integrate frequently). The key goals of CI are to find and address bugs more quickly, improve software quality, and reduce the time it takes to validate and release new software updates.

Continuous integration focuses on smaller commits and smaller code changes to integrate. A developer commits code at regular intervals, at minimum once a day. The developer pulls code from the code repository to ensure the code on the local host is merged before pushing to the build server. At this stage the build server runs the various tests and either accepts or rejects the code commit.

The basic challenges of implementing CI include more frequent commits to the common codebase, maintaining a single source code repository, automating builds, and automating testing. Additional challenges include testing in similar environments to production, providing visibility of the process to the team, and allowing developers to easily obtain any version of the application.

## Continuous Delivery (CD)

Continuous delivery (CD) is a software development practice where code changes are automatically built, tested, and prepared for production release. It expands on continuous integration by deploying all code changes to a testing environment, a production environment, or both after the build stage has been completed. Continuous delivery can be fully automated with a workflow process or partially automated with manual steps at critical points. When continuous delivery is properly implemented, developers always have a deployment-ready build artifact that has passed through a standardized test process. With continuous deployment, revisions are deployed to a production environment.

CD can help you discover and address bugs early in the delivery process before they grow into larger problems later. Our team can easily perform additional types of code tests because the entire process has been automated. With the discipline of more testing more frequently, we can iterate faster with immediate feedback on the impact of changes. This enables the team to drive quality code with a high assurance of stability and security. Developers will know through immediate feedback whether the new code works and whether any breaking changes or bugs were introduced. Mistakes caught early on in the development process are the easiest to fix.

## Purpose of the document

This section discusses the ways in which you can use the best practices when implementing a CI/CD model and how we used these practices into our project.

Continuous integration, delivery and deployment are software development practices born out of the DevOps movement. They make the process of building, testing and releasing software more efficient and get working software into the hands of users more quickly than traditional methods. Done well, a CI/CD pipeline enables teams to deliver working software at pace and get timely feedback on their latest changes.

Each commit triggers a set of automated tests to provide prompt feedback on the change. In order to reap the benefits of continuous integration, it's essential for everyone to share their changes with the rest of your team by pushing to main (master) and to update their working copy in order to receive everyone else's changes. As a general rule aim to commit to main (master) at least once a day.

Security first approach, Strong privilege management is a necessity in fast-moving application pipelines. Integration with secrets management tools and a granular IAM policy engine like AWS IAM is crucial, along with integrations with directory services. Security teams helps to define the appropriate least privilege access models needed for all stages of application development and deployment, and then implement this in a centralized tool/service whenever possible.

## Scope of the document

In this section we want to highlight the key points we considered when developing this system and how we achieve each one of them:

**Highly Availability** is the ability of a system to operate continuously without failing for a designated period. Availability can be measured relative to a system being 100% operational or never failing -- meaning it has no outages. In order to guarantee the high availability for our system we created to elastic containers, this way when we are running a updated in one of the containers the other will still be working, not affecting the end-user usability.

**Scalable:** A scalable cloud architecture is made possible through virtualization. Different then a physical machine where resources and performance are relatively set, virtual machines (VMs) are highly flexible and can be easily scaled up or down. They can be moved to a different server or hosted on multiple servers at once; workloads and applications can be shifted to larger VMs as needed. When using CodePipeline, each pipeline run executes independently and in parallel with

each other. Since CodePipeline is a managed service, we can run many pipelines in parallel without having to deal with additional infrastructure.

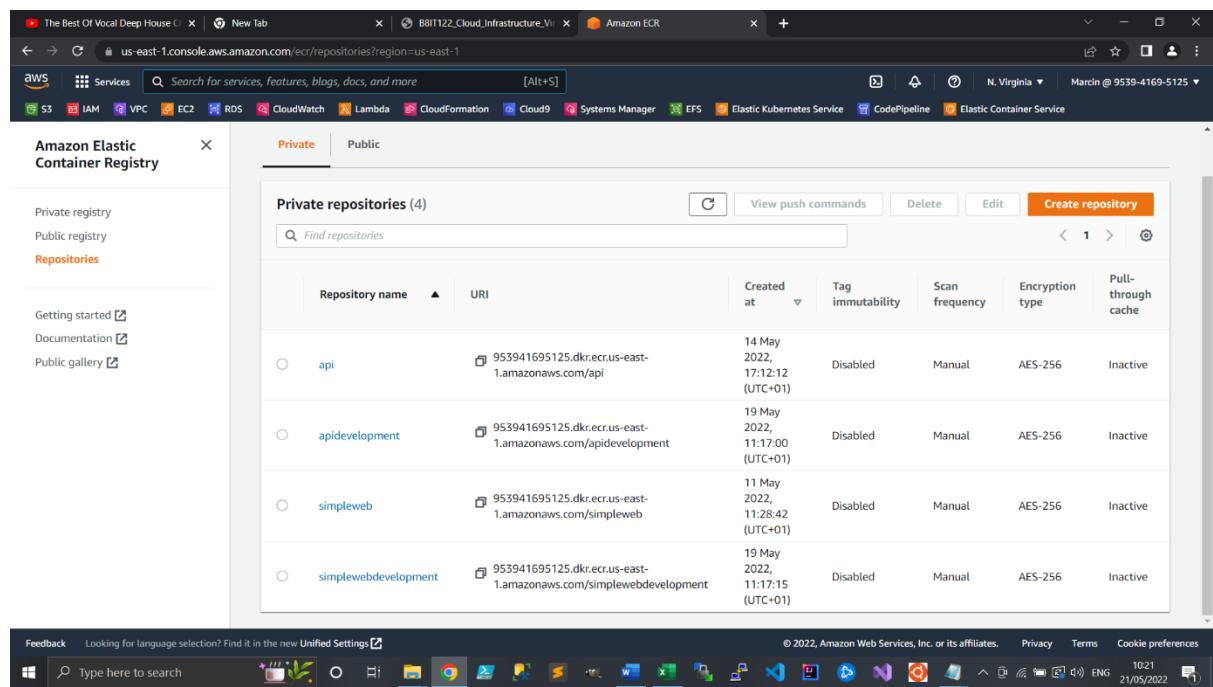
For this project we used two different Availability Zones, in distinct locations within an AWS Region. They provide inexpensive, low-latency network connectivity to other Availability Zones in the same AWS Region.

**and Secure System:** Cluster is deployed in private subnet 1 and private subnet 2, DB is receiving traffic just from security group of hypervisors and Cloud9 instance. Containers receive traffic just from load balancer security group and DB security group. Load balancer security group accept traffic just from ports that are used by project. Nothing else is open.

## Environment information

Below you'll find the tools we used to achieve the requirements:

**ECR – Amazon Elastic Container Registry** - Containers Repository to hold our containers that are built by CodeBuild. ECR is a fully managed container registry offering high-performance hosting, so you can reliably deploy application images and artifacts anywhere. We have 4 container repositories created for project:



The screenshot shows the AWS ECR console with four private repositories listed:

Repository name	URI	Created at	Tag immutability	Scan frequency	Encryption type	Pull-through cache
api	953941695125.dkr.ecr.us-east-1.amazonaws.com/api	14 May 2022, 17:12:12 (UTC+01)	Disabled	Manual	AES-256	Inactive
apidevelopment	953941695125.dkr.ecr.us-east-1.amazonaws.com/apidevelopment	19 May 2022, 11:17:00 (UTC+01)	Disabled	Manual	AES-256	Inactive
simpleweb	953941695125.dkr.ecr.us-east-1.amazonaws.com/simpleweb	11 May 2022, 11:28:42 (UTC+01)	Disabled	Manual	AES-256	Inactive
simplewebdevelopment	953941695125.dkr.ecr.us-east-1.amazonaws.com/simplewebdevelopment	19 May 2022, 11:17:15 (UTC+01)	Disabled	Manual	AES-256	Inactive

**ECS – Amazon Elastic Container Service (Amazon ECS)** is a highly scalable and fast container orchestration service. We used it to run, stop, and manage containers on a cluster. With Amazon ECS, our containers are defined in a task definition that we use to run an individual task or task within a service. In this context, a service is a configuration that you can use to run and maintain a specified number of tasks simultaneously in a cluster. We can run your tasks and services on a

serverless infrastructure that's managed by AWS Fargate. Alternatively, for more control over your infrastructure, we can run your tasks and services on a cluster of Amazon EC2 instances that we manage. For propose of this project we run ECS on EC2.

Cluster : dbscluster

Cluster ARN: arn:aws:ecs:us-east-1:953941695125:cluster/dbscluster

Status: ACTIVE

Registered container instances: 3

Pending tasks count	0 Fargate, 0 EC2, 0 External
Running tasks count	0 Fargate, 6 EC2, 0 External
Active service count	0 Fargate, 4 EC2, 0 External
Draining service count	0 Fargate, 0 EC2, 0 External

Services Tasks ECS Instances Metrics Scheduled Tasks Tags Capacity Providers

Feedback Looking for language selection? Find it in the new Unified Settings

**CodeCommit** – AWS Code Repository. CodeCommit is a secure, managed source control service that hosts our repositories. We have created 2 repositories with 2 branches each main and development:

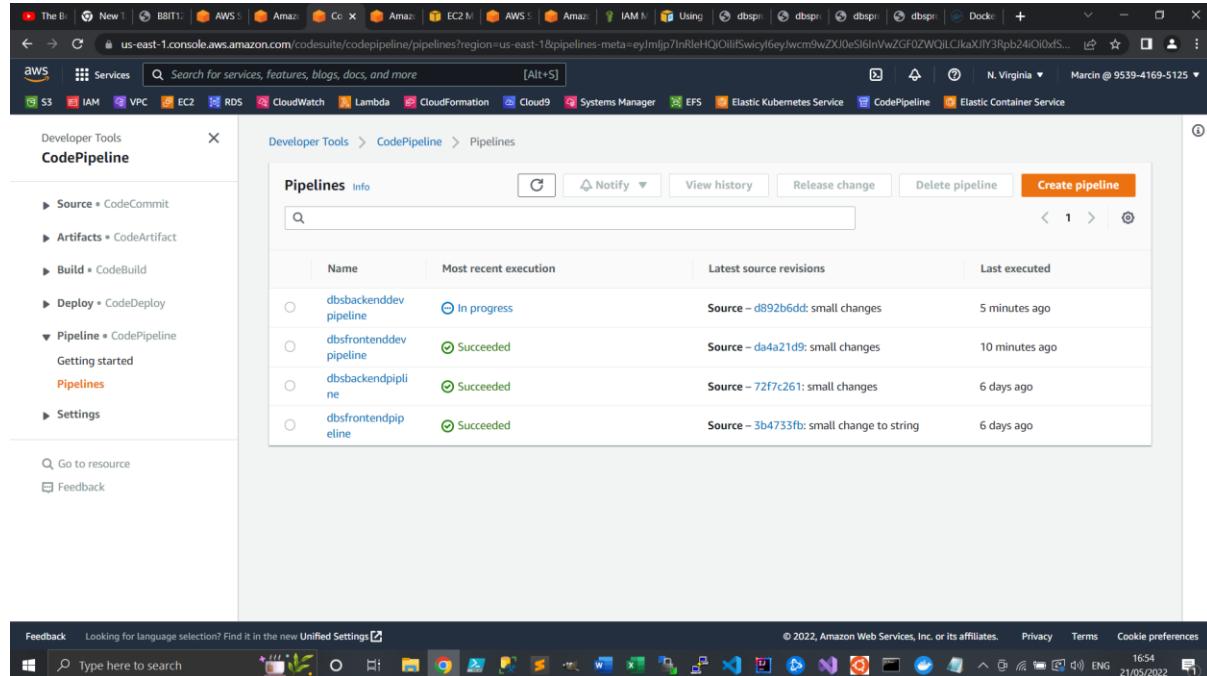
Developer Tools > CodeCommit > Repositories

Repositories

Name	Description	Last modified	Clone URL
api	-	1 day ago	HTTPS SSH HTTPS (GRC)
simpleweb	-	1 day ago	HTTPS SSH HTTPS (GRC)

Feedback Looking for language selection? Find it in the new Unified Settings

**CodePipeline** – AWS Service to Build CI/CD. Fully managed continuous delivery service that helps you automate your release pipelines for fast and reliable application and infrastructure updates. We build 4 pipelines: frontend, backend, frontend development and backend development.

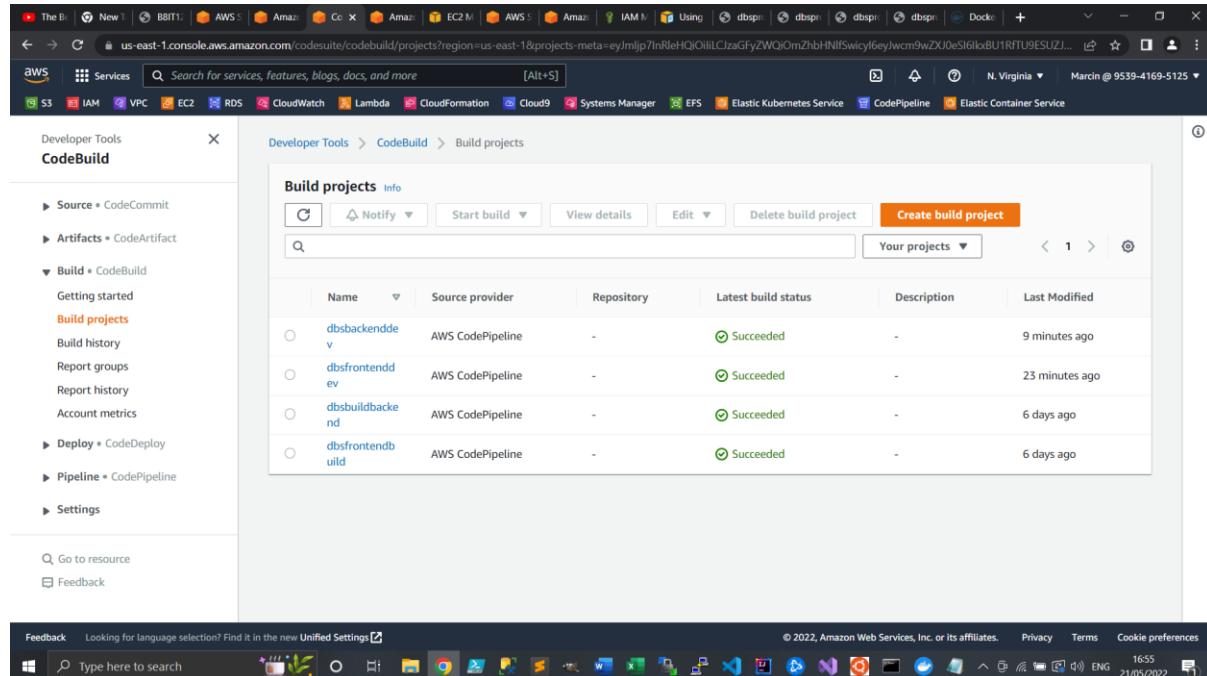


The screenshot shows the AWS CodePipeline console with the following details:

- Pipelines** table:
 

Name	Most recent execution	Latest source revisions	Last executed
dbsbackenddev pipeline	In progress	Source - d892b6dd: small changes	5 minutes ago
dbsfrontenddev pipeline	Succeeded	Source - da4a21d9: small changes	10 minutes ago
dbsbackendpipeline	Succeeded	Source - 72f7c261: small changes	6 days ago
dbsfrontendpipeline	Succeeded	Source - 3b4733fb: small change to string	6 days ago

**CodeBuild** – AWS Service that compiles source code runs tests and produces software packages to deploy. For our pipeline we have 4 build projects:

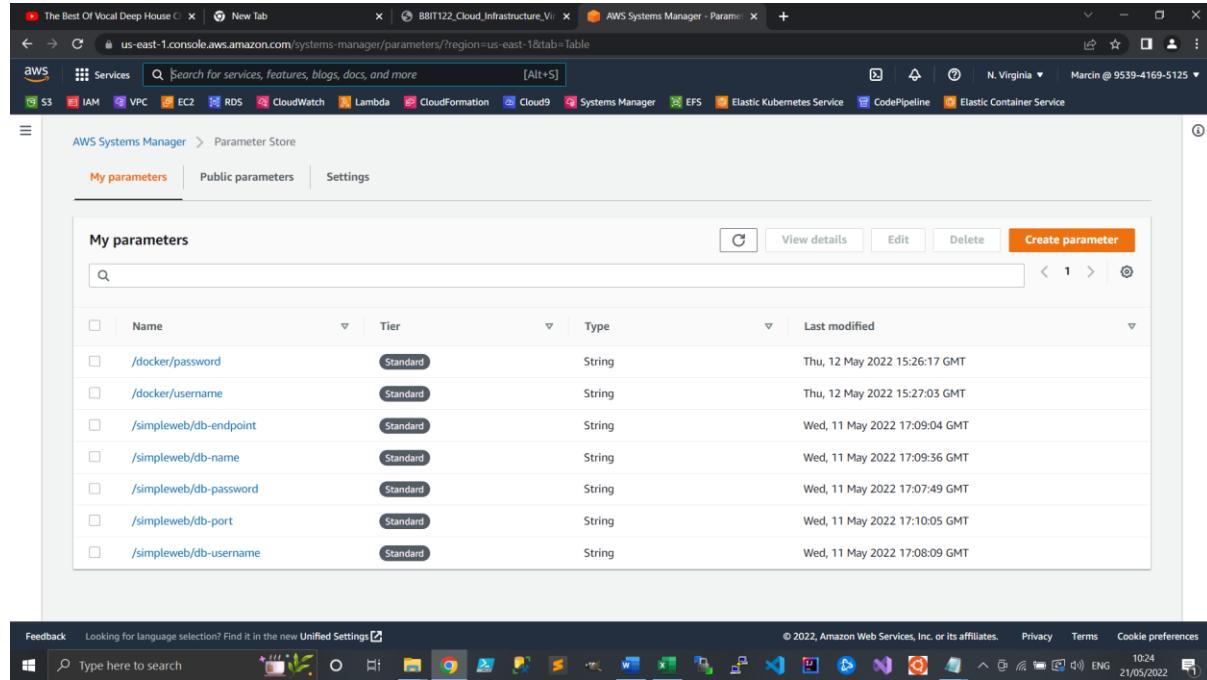


The screenshot shows the AWS CodeBuild console with the following details:

- Build projects** table:
 

Name	Source provider	Repository	Latest build status	Description	Last Modified
dbsbackenddev	AWS CodePipeline	-	Succeeded	-	9 minutes ago
dbsfrontenddev	AWS CodePipeline	-	Succeeded	-	23 minutes ago
dbsbuildbackend	AWS CodePipeline	-	Succeeded	-	6 days ago
dbsfrontedbuid	AWS CodePipeline	-	Succeeded	-	6 days ago

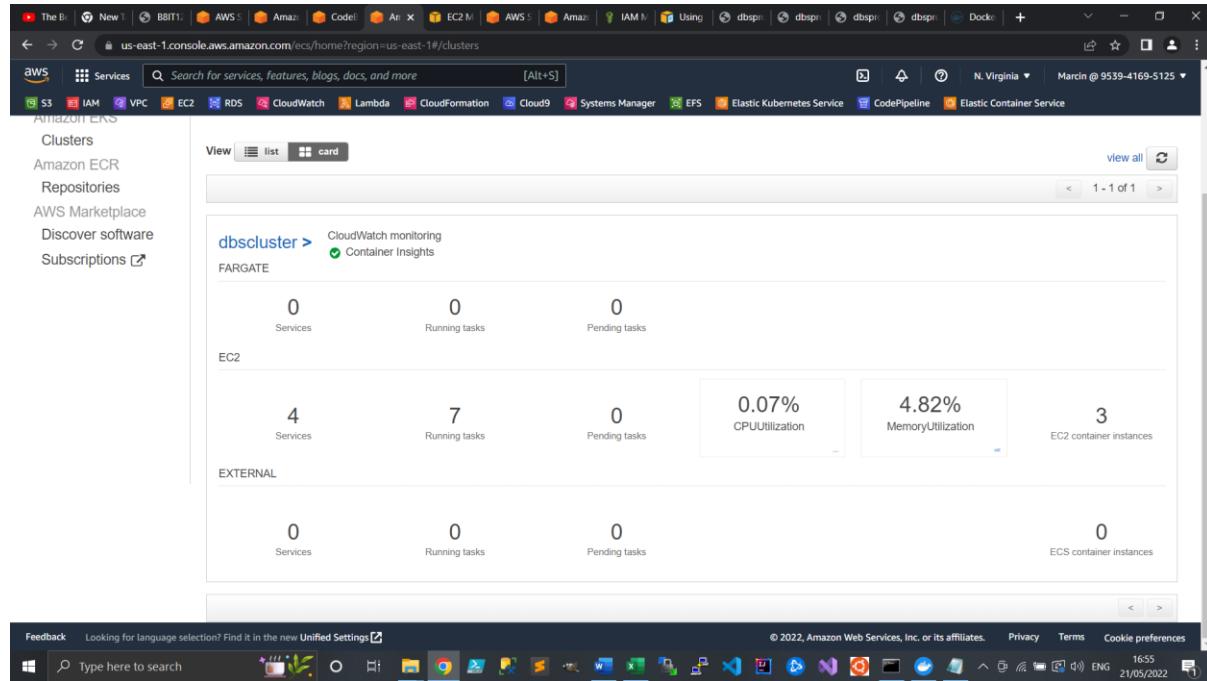
**Parameter store** – AWS System Manager parameter store for holding secrets. We used this service to store data such as passwords, database strings, Amazon Machine Image (AMI) IDs, and license codes as parameter values.



The screenshot shows the AWS Systems Manager Parameter Store interface. The top navigation bar includes tabs for 'My parameters', 'Public parameters', and 'Settings'. A search bar is at the top. Below is a table of parameters:

Name	Tier	Type	Last modified
/docker/password	Standard	String	Thu, 12 May 2022 15:26:17 GMT
/docker/username	Standard	String	Thu, 12 May 2022 15:27:03 GMT
/simpleweb/db-endpoint	Standard	String	Wed, 11 May 2022 17:09:04 GMT
/simpleweb/db-name	Standard	String	Wed, 11 May 2022 17:09:36 GMT
/simpleweb/db-password	Standard	String	Wed, 11 May 2022 17:07:49 GMT
/simpleweb/db-port	Standard	String	Wed, 11 May 2022 17:10:05 GMT
/simpleweb/db-username	Standard	String	Wed, 11 May 2022 17:08:09 GMT

**CloudWatch event** – AWS CloudWatch logs and events. This service describes changes in Amazon Web Services (AWS) resources, also provide for logs.



The screenshot shows the AWS CloudWatch Events interface. On the left, there's a sidebar with links for 'Clusters', 'Amazon ECR', 'Repositories', 'AWS Marketplace', 'Discover software', and 'Subscriptions'. The main area displays a summary of tasks and resources:

Category	Value
Services	0
Running tasks	0
Pending tasks	0
CPUUtilization	0.07%
MemoryUtilization	4.82%
EC2 container instances	3
ECS container instances	0

**Cloud9** – AWS Cloud IDE. Cloud-based integrated development environment (IDE) that users use to write, run, and debug code or connect to DB

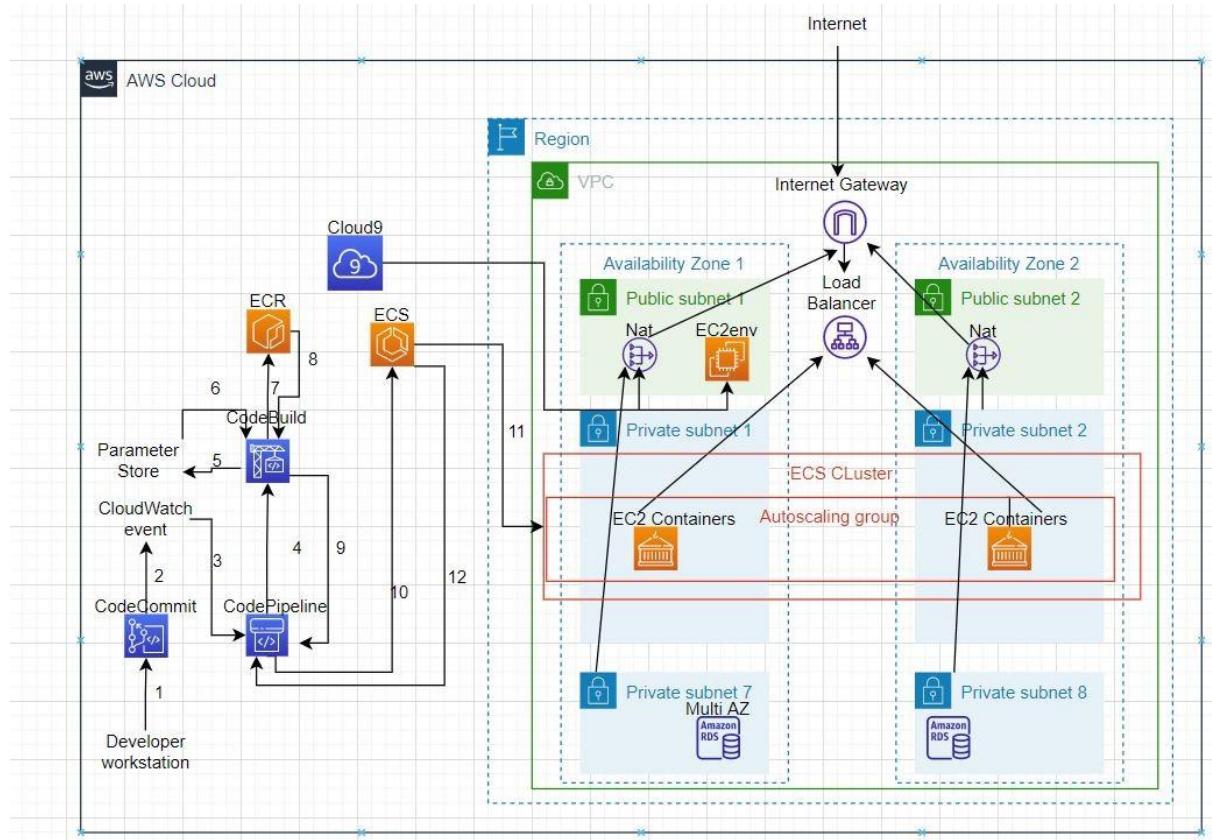
## RDS – AWS Service for DB

The screenshot shows the AWS RDS console with the following details:

- Region:** us-east-1
- Databases:** simplewebdbprivate (MySQL Community, db.t2.micro, Available)
- Subnet groups:** simplewebdbprivate
- Parameter groups:** simplewebdbprivate
- Option groups:** simplewebdbprivate
- Custom engine versions:** simplewebdbprivate
- Events:** simplewebdbprivate
- Event subscriptions:** simplewebdbprivate

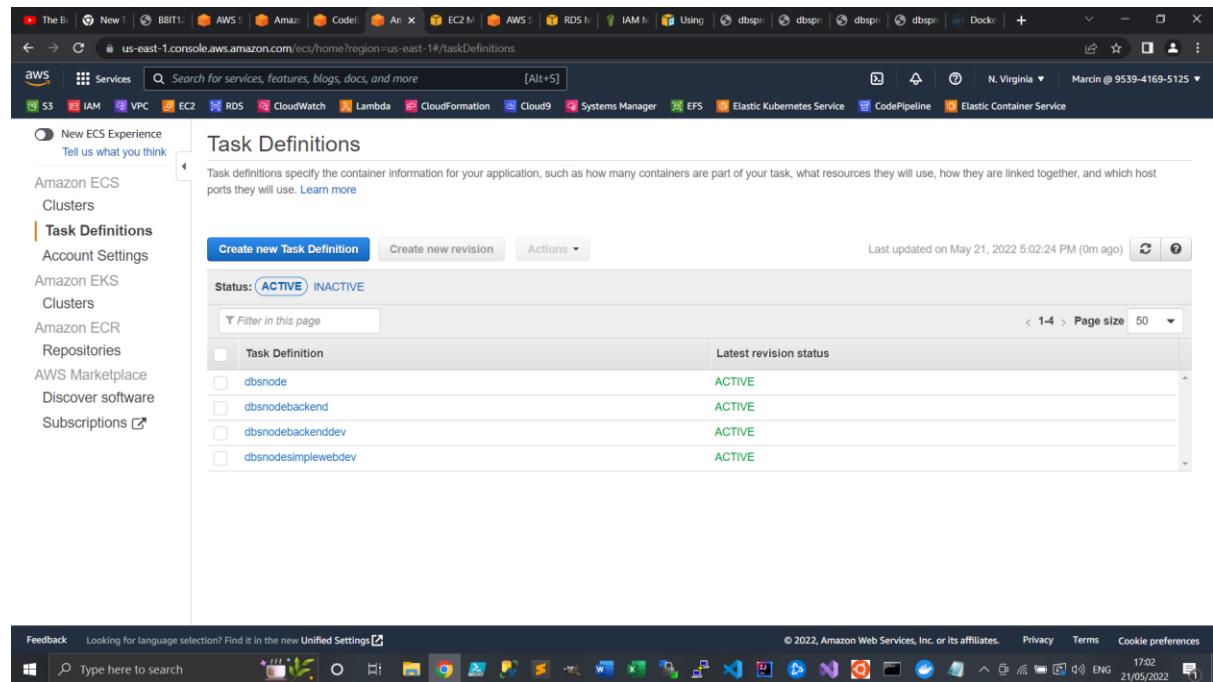
## Architecture

### Design Specification



Our ECS containers cluster is deployed on Private subnet 1 and Private subnet 2 on EC2's. Containers accept traffic just from security groups of load balancer and DB MSQl which is on AWS RDS. Containers outbound traffic is going through NAT which is different for Private subnet 1 and Private subnet 2. Autoscaling group deploy 3 EC2 servers for the cluster as desired but up to 5 if needed by ECS cluster which check if CPU usage is above 80% on already running servers. ECS cluster have services which maintain state of containers by running task definition. Service which maintains state of frontend production have min of 2 task running (2 containers) also backend production service have min 2 task running and max 5. Services which maintain frontend and backend development have min 1 task running and max 2. Deployment is set as spread AZ for production and binpack for development. For propose of this project we have created 4 different task definitions and 4 services:

## Task Definitions



The screenshot shows the AWS ECS Task Definitions page. The URL is [us-east-1.console.aws.amazon.com/ecs/home?region=us-east-1#/taskDefinitions](https://us-east-1.console.aws.amazon.com/ecs/home?region=us-east-1#/taskDefinitions). The page title is "Task Definitions". A sidebar on the left lists navigation options: "New ECS Experience", "Tell us what you think", "Amazon ECS Clusters", "Task Definitions" (which is selected and highlighted in orange), "Account Settings", "Amazon EKS Clusters", "Amazon ECR Repositories", "AWS Marketplace Discover software", and "Subscriptions". The main content area has a table with the following data:

Task Definition	Latest revision status
dbsnode	ACTIVE
dbsnodebackend	ACTIVE
dbsnodebackenddev	ACTIVE
dbsnodesimplewebdev	ACTIVE

At the bottom of the page, there is a footer with links to "Feedback", "Language selection", "Privacy", "Terms", and "Cookie preferences". The footer also shows the date "21/05/2022" and the time "17:02".

## Services

Clusters

Task Definitions

Account Settings

Amazon EKS

Clusters

Amazon ECR

Repositories

AWS Marketplace

Discover software

Subscriptions

Cluster ARN: arn:aws:ecs:us-east-1:953941695125:cluster/dbscluster

Status: ACTIVE

Registered container instances: 3

Pending tasks count: 0 Fargate, 0 EC2, 0 External

Running tasks count: 0 Fargate, 6 EC2, 0 External

Active service count: 0 Fargate, 4 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

Last updated on May 21, 2022 5:03:22 PM (0m ago)

Service Name	Status	Service t...	Task Definition	Desired t...	Running t...	Launch t...	Platform ...
dbsbackendnodeservice	ACTIVE	REPLICA	dbsnodebackend:3	2	2	EC2	—
dbsfrontendnodeservice	ACTIVE	REPLICA	dbsnode:15	2	2	EC2	—
dbsbackendnodeservice	ACTIVE	REPLICA	dbsnodebackenddev:2	1	1	EC2	—
dbsfrontendnodeservice	ACTIVE	REPLICA	dbsnodesimplewebdev:2	1	1	EC2	—

Load balancer accept traffic just needed by our apps:

Feedback Looking for language selection? Find it in the new Unified Settings

Services

Search for services, features, blogs, docs, and more

Resered Instances New

Dedicated Hosts

Scheduled Instances

Capacity Reservations

Images

AMIs New

AMI Catalog

Elastic Block Store

Volumes New

Snapshots New

Lifecycle Manager New

Network & Security

Security Groups

Elastic IPs

Placement Groups

Key Pairs

Network Interfaces

Load Balancing

Load Balancers

Security group name: LB SG

Security group ID: sg-0745380c11684d2c4

Description: Allow traffic for LB

VPC ID: vpc-0f47e6a0ab92b6204

Owner: 953941695125

Inbound rules count: 4 Permission entries

Outbound rules count: 1 Permission entry

Inbound rules Outbound rules Tags

You can now check network connectivity with Reachability Analyzer

Run Reachability Analyzer

Inbound rules (4)

Filter security group rules

rule...	IP version	Type	Protocol	Port range	Source	Description
9269c63	IPv4	Custom TCP	TCP	7000	0.0.0.0/0	—
dae470b	IPv4	Custom TCP	TCP	8000	0.0.0.0/0	—
ff7270bc	IPv4	HTTP	TCP	80	0.0.0.0/0	—
2de6e1	IPv4	Custom TCP	TCP	5000	0.0.0.0/0	—

Have 4 listeners and 4 target groups:

```

import os
import json
import pymysql

def lambda_handler(event, context):
    # Get the connection string from the environment variables
    connection_string = os.environ['DB_CONNECTION_STRING']
    connection = pymysql.connect(**connection_string)
    cursor = connection.cursor()
    cursor.execute("SELECT * FROM users")
    rows = cursor.fetchall()
    response = {
        'statusCode': 200,
        'body': json.dumps(rows)
    }
    return response

```

Cloud9 Instance accept traffic just from DB security group and needed by Cloud9 service, users can easily access DB from IDE.

CI/CD steps explained:

2 pipelines were developed: Frontend and Backend, both using the tools listed below:

1. Developer ran git push command to CodeCommit successfully
2. CodeCommit received update successfully
3. CloudWatch see event and is sending information to CodePipline
4. CodePipeline start pipeline for which event was triggered, running test on source which is CodeCommit if results are positive going to next step
5. CodeBuild project is starting building docker image, service is gathering necessary information from parameter store
6. CodeBuild Project received necessary details from parameter store and is continue building docker image
7. Once container is ready. CodeBuild is sending docker image to ECR repository
8. uploaded to ECR is successful
9. CodeBuild is informing pipeline of project completion successfully
10. CodePipeline pipeline is sending information for deployment to ECS

11. ECS is sending information do cluster to start deployment
12. ECS is sending information to pipeline of deployment completion

## Access Procedure

In order to have a more secure environment, group role was created:

For the developers group role was created: Giving full access to CodeCommit repository, Read-only access to CodePipeline and access to Cloud9 as user from where they can connect to database

The screenshot shows the AWS IAM Management Console. On the left, a sidebar navigation includes 'Identity and Access Management (IAM)', 'Access management' (with 'User groups' selected), and 'Access reports'. The main content area is titled 'Summary' for a user group named 'Developersdbspj'. It shows the group was created on May 13, 2022, at 20:51 (UTC+01:00). The 'Permissions' tab is selected, showing a table of managed policies:

Policy name	Type	Description
AWSCodeCommitFullAccess	AWS managed	Provides full access to AWS CodeCommit via the AWS Management Console
AWSCodePipelineReadOnlyAccess	AWS managed	Provides read only access to AWS CodePipeline via the AWS Management Console
AWSCloud9User	AWS managed	Provides permission to create AWS Cloud9 development environments and to

### IAM Groups

IAM Groups are a way of grouping IAM users and IAM roles. Permissions given to an IAM Group are passed onto their group members (users and roles).

The IAM permissions given to an IAM Group (or IAM user or IAM role) determine which AWS API commands can be executed using the AWS CLI or any of the many AWS SDKs.

### Resources:

<https://www.ukessays.com/essays/computer-science/ci-and-cd-in-aws-code-pipeline.php>

<https://docs.aws.amazon.com/whitepapers/latest/practicing-continuous-integration-continuous-delivery/practicing-continuous-integration-continuous-delivery.pdf>

<https://aws.amazon.com/ecr/>

[https://mng.workshop.aws/ssm/capability\\_hands-on\\_labs/parameterstore.html](https://mng.workshop.aws/ssm/capability_hands-on_labs/parameterstore.html)

<https://www.techtarget.com/searchdatacenter/definition/high-availability>

<https://www.vmware.com/topics/glossary/content/cloud-scalability.html#:~:text=A%20scalable%20cloud%20architecture%20is,easily%20scaled%20up%20or%20down.>

<https://www.jetbrains.com/teamcity/ci-cd-guide/ci-cd-best-practices/>