

ConfigurationManual

MScResearchProject
MScCloud Computing

Bramha Theja Gadikota
StudentID:23197994

SchoolofComputing
NationalCollegeofIreland

Supervisor: Prof.Shaguna gupta

National College of Ireland
Project Submission Sheet School of
Computing



Student Name:	Bramha Theja Gadikota
Student ID:	23197994
Programme:	Cloud Computing
Year:	2024
Module:	MSc Research Project
Supervisor:	Prof. Shaguna Gupta
Submission Due Date:	12/12/2023
Project Title:	Configuration Manual
Word Count:	710
Page Count:	10

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:	Bramha Theja Gadikota
Date:	10th December 2024

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	☐
Attach a Moodle submission receipt of the online project submission , to each project (including multiple copies).	☐
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.	☐

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):	

Configuration Manual

Bramha Theja Gadikota
23197994

INTRODUCTION

This configuration manual provides in-depth, step-by-step instructions for installing, configuring, and deploying all of the software, tools, and files that are necessary for the implementation of the proposed system.

PREREQUISITES

- **Programming:** Python, Flask
- **Containerization:** Docker
- **Cloud Services:** Azure ACR, AKS, AWS, ECR, EKS
- **Tools:** VS Code, Locust, Kubernetes (kubectl)

WORKFLOW

1. Application Development:

- Implement a Flask web application with endpoints for:
 - ▢ **Home page (/):** Serves the main web interface.
 - ▢ **Start CPU Load (/start-load):** Initiates CPU stress testing with optional duration and core configuration.
 - ▢ **Stop CPU Load (/stop-load):** Stops ongoing stress testing.
 - ▢ **Status (/status):** Checks if CPU stress testing is running.
- Use threading and multiprocessing for parallel stress simulation across all CPU cores.

2. Containerization:

- Create a Dockerfile to package the Flask application into a container image.
 - ▢ Include necessary dependencies, Python environment, and application files.
 - ▢ Expose the application on port 5000.
- Build and test the Docker image locally.

3. Azure Deployment:

- **Azure Container Registry (ACR):**
 - ▢ Push the Docker image to ACR for centralized image management.
- **Azure Kubernetes Service (AKS):**
 - ▢ Deploy the application to an AKS cluster.

4. AWS Deployment:

- **Elastic Container Registry (ECR):**
 - ▢ Push the Docker image to ECR for centralized image management.
- **Elastic Kubernetes Service (EKS):**
 - ▢ Deploy the application to an EKS cluster.

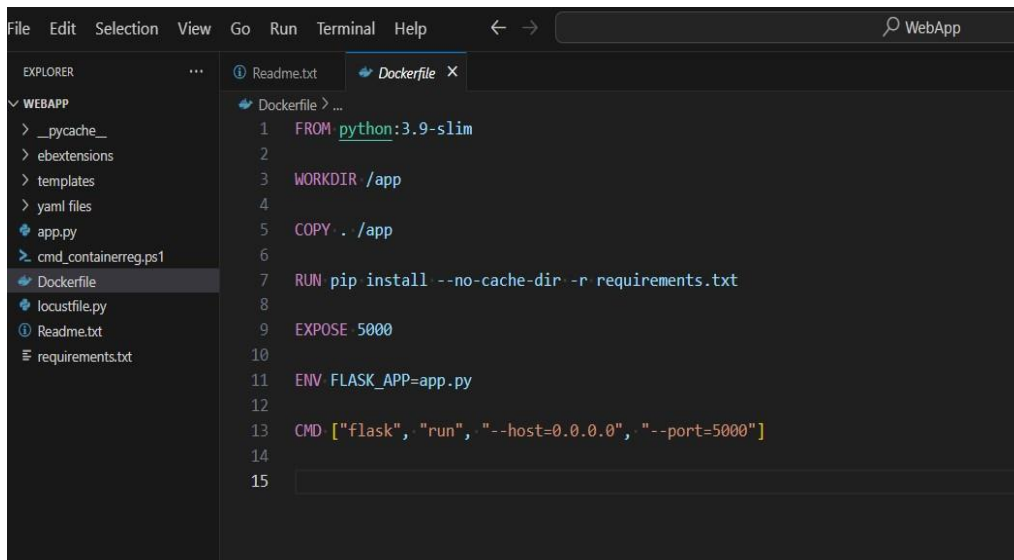
5. Locust:

- Load Test

IMPLEMENTATION

CONTAINERIZATION

Dockerfile:



Docker Build Commands:

Build Docker Image:

```
docker build -t cpuloadgenflask .
```

Run Docker Locally: `docker run -p`

```
5000:5000 -d cpuloadgenflask flask run
```

AZURE DEPLOYMENT

Azure Container Registry (ACR) Login

to ACR:

```
az acr login --name cpuloadgenacr
```

Push Docker Image: `docker tag cpuloadgenflask`

```
cpuloadgenacr.azurecr.io/cpuloadgenflask docker push
```


```
cpuloadgenacr.azurecr.io/cpuloadgenflask
```

Azure Kubernetes Service (AKS) Connect

to Cluster:

```
az aks get-credentials --resource-group cpuloadgenrg --name cpuloadgenaks --  
overwriteexisting
```

Deploy to AKS:



```
! deployment.yaml X  
! deployment.yaml > apiVersion  
1  apiVersion: apps/v1  
2  kind: Deployment  
3  metadata:  
4    name: flask-deployment  
5  spec:  
6    replicas: 3  
7    selector:  
8      matchLabels:  
9        app: flask-app  
10   template:  
11     metadata:  
12       labels:  
13         app: flask-app  
14     spec:  
15       containers:  
16         - name: flask-container  
17           image: cpuloadgenacr.azurecr.io/cpuloadgenflask  
18           ports:  
19             - containerPort: 5000
```

```
kubectl apply -f deployment.yaml
```

```
! service.yaml X
! service.yaml > {} spec > {} selector > app
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: flask-service
5  spec:
6    type: LoadBalancer
7    ports:
8      - port: 80
9        targetPort: 5000
10   selector:
11     app: flask-app
```

```
kubectl apply -f service.yaml
```

```
kubectl get service flask-service
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
flask-service	NodePort	10.100.237.2	<none>	80:30007/TCP	14d

AWS DEPLOYMENT

Elastic Container Registry (ECR) Login

to ECR:

```
aws ecr get-login-password --region eu-west-2 | docker login --username AWS --password-stdin 820242902892.dkr.ecr.eu-west-2.amazonaws.com
```

Push Docker Image:

```
docker tag cpuloadgenflask:latest 820242902892.dkr.ecr.eu-west-2.amazonaws.com/cpuloadgenecr:latest
```

```
docker push 820242902892.dkr.ecr.eu-west-2.amazonaws.com/cpuloadgenecr:latest
```

Elastic Kubernetes Service (EKS) Create

EKS Cluster:

```
eksctl create cluster --name cpuloadgeneks-cluster --version 1.31 --region eu-west-2 --nodegroup-name linux-nodes1 --node-type t3.medium --nodes 3 --nodes-min 1 --nodes-max 10 --managed
```

Connect to Cluster:

```
aws eks update-kubeconfig --name cpuloadgeneks-cluster
```

Deploy to EKS:


```

! deployment.yaml X
! deployment.yaml > apiVersion
1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: flask-deployment1
5  spec:
6    replicas: 1
7    selector:
8      matchLabels:
9        app: flask-app
10   template:
11     metadata:
12       name: myapp-pod
13     labels:
14       app: flask-app
15     spec:
16       containers:
17         name: flask-container
18         image: 820242902892.dkr.ecr.eu-west-2.amazonaws.com/cpuloadgenecr:latest
19         ports:
20         containerPort: 5000

```

kubectl apply -f deployment.yaml --validate=false

```

! service.yaml X
! service.yaml > apiVersion
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: flask-service1
5  spec:
6    type: LoadBalancer
7    ports:
8      - port: 80
9        targetPort: 5000
10   selector:
11     app: flask-app

```

kubectl apply -f service.yaml --validate=false

kubectl get service flask-service1

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
flask-service1	LoadBalancer	10.100.211.170	a14e82085a0ac4ba5af09fb49438ab86-1898328803.eu-west-2.elb.amazonaws.com	80:32015/TCP	6d6h

INSTALLATION OF MONITORING AND VISUALISATION TOOLS

Created a kubernetes service of Grafana and Prometheus using YAML file. **GRAFANA**

```
! grafana_service1.yaml  ! grafana_deployment1.yaml X
yaml files > ! grafana_deployment1.yaml > {} spec > {} template > {} spec > [ ] volumes > {} 0
1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    - name: grafana1
5    - namespace: default
6    - labels:
7      - app: grafana
8  spec:
9    - replicas: 1
10   - selector:
11     - matchLabels:
12       - app: grafana
13   - template:
14     - metadata:
15       - labels:
16         - app: grafana
17     - spec:
18       - containers:
19         - name: grafana
20         - image: grafana/grafana:latest
21         - ports:
22           - containerPort: 3000
23         - volumeMounts:
24           - name: grafana-storage
25             mountPath: /var/lib/grafana
26         - volumes:
27           - name: grafana-storage
28             emptyDir: {}
29
```

kubectl apply -f grafana_deployment1.yaml

```
! grafana_service1.yaml X ! grafana_deployment1.yaml
yaml files > ! grafana_service1.yaml > apiVersion
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: grafana1
5    namespace: default
6    labels:
7      app: grafana
8  spec:
9    selector:
10     app: grafana
11    ports:
12     - protocol: TCP
13       port: 80
14       targetPort: 3000
15     type: LoadBalancer
16
```

kubectl apply -f grafana_service1.yaml

kubectl get service grafana1

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
grafana1	LoadBalancer	10.100.26.196	a1cd4b035250a42d1a98cdd9ef46ebfe-1636237739.eu-west-2.elb.amazonaws.com	80:31949/TCP	6d5h

PROMETHEUS

```

! prometheus_deployment1.yaml X ! prometheus_service1.yaml
yaml files > ! prometheus_deployment1.yaml > {} spec > {} template > {} spec > {} volumes > {} 1
1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: prometheus
5    namespace: default
6    labels:
7      app: prometheus
8  spec:
9    replicas: 1
10   selector:
11     matchLabels:
12       app: prometheus
13   template:
14     metadata:
15       labels:
16         app: prometheus
17     spec:
18       containers:
19         - name: prometheus
20           image: prom/prometheus:latest
21           args:
22             - "--config.file=/etc/prometheus/prometheus.yml"
23             - "--storage.tsdb.path=/prometheus"
24             - "--web.enable-lifecycle"
25           ports:
26             - containerPort: 9090
27           volumeMounts:
28             - name: prometheus-config-volume
29               mountPath: /etc/prometheus
30             - name: prometheus-data
31               mountPath: /prometheus
32           volumes:
33             - name: prometheus-config-volume
34               configMap:
35                 name: prometheus-config
36             - name: prometheus-data
37               emptyDir: {}
38

```

kubectl apply -f prometheus_deployment1.yaml

```

! prometheus_deployment1.yaml ! prometheus_service1.yaml X
yaml files > ! prometheus_service1.yaml > apiVersion
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: prometheus
5    namespace: default
6    labels:
7      app: prometheus
8  spec:
9    selector:
10     app: prometheus
11   ports:
12     - protocol: TCP
13       port: 80
14       targetPort: 9090
15   type: LoadBalancer
16

```

kubectl apply -f prometheus_service1.yaml kubectl get service prometheus

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
prometheus	LoadBalancer	10.100.180.139	a828207e0a9474c66872e35ba98d30e2-1355634064.eu-west-2.elb.amazonaws.com	80:31082/TCP	6d5h

GET ALL RUNNING SERVICE

```
kubectl get svc -n default
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
flask-service	NodePort	10.100.237.2	<none>	80:30007/TCP	15d
flask-service1	LoadBalancer	10.100.211.170	a14e82085a0ac4ba5af09fb49438ab86-1898328803.eu-west-2.elb.amazonaws.com	80:32015/TCP	6d6h
grafana	NodePort	10.100.62.97	<none>	80:32000/TCP	6d6h
grafana1	LoadBalancer	10.100.26.196	a1cd4b035250a42d1a98cdd9ef46ebfe-1636237739.eu-west-2.elb.amazonaws.com	80:31949/TCP	6d5h
kubernetes	ClusterIP	10.100.0.1	<none>	443/TCP	15d
prometheus	LoadBalancer	10.100.180.139	a828207e0a9474c66872e35ba98d30e2-1355634064.eu-west-2.elb.amazonaws.com	80:31082/TCP	6d5h

LINEAR MODEL FOR CPU RESOURCE SCHEDULING

```
24
25 def linear_cpu_prediction():
26     """#linear model for cpu prediction
27     """
28     from sklearn.linear_model import LinearRegression
29     from sklearn.metrics import mean_squared_error
30
31     """# Train the model
32     """
33     model = LinearRegression()
34     model.fit(X_train, y_train)
35
36     """# Evaluate the model
37     """
38     y_pred = model.predict(X_test)
39     mse = mean_squared_error(y_test, y_pred)
40     print(f"Mean Squared Error: {mse}")
41     return y_pred
42
43
44 def schedule_pod():
45     """def schedule_pod_resource():
46     """
47     return linear_cpu_prediction()
48     predicted_cpu, predicted_memory = schedule_pod_resource()
49
50     print(f"Predicted CPU Usage: {predicted_cpu}m")
51     print(f"Predicted Memory Usage: {predicted_memory}Mi")
52
53     pod = client.V1Pod(
54         metadata=client.V1ObjectMeta(name="novelty-pod"),
55         spec=client.V1PodSpec(
56             containers=[
57                 client.V1Container(
58                     name="novelty-container",
59                     image="nginx",
60                     resources=client.V1ResourceRequirements(
61                         requests={
62                             "cpu": f"{predicted_cpu}m",
63                             "memory": f"{predicted_memory}Mi"
64                         }
65                     )
66                 )
67             ]
68         )
69     )
```