

# Configuration Manual

MSc Research Project  
Cloud Computing

Anant Sakharam Pednekar  
Student ID: 21188947

School of Computing  
National College of Ireland

Supervisor: Rejwanul Haque

National College of Ireland  
Project Submission Sheet  
School of Computing



<b>Student Name:</b>	Anant Sakharam Pednekar
<b>Student ID:</b>	21188947
<b>Programme:</b>	Cloud Computing
<b>Year:</b>	2023
<b>Module:</b>	MSc Research Project
<b>Supervisor:</b>	Rejwanul Haque
<b>Submission Due Date:</b>	15/12/2023
<b>Project Title:</b>	Configuration Manual
<b>Word Count:</b>	XXX
<b>Page Count:</b>	5

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.

<b>Signature:</b>	Anant
<b>Date:</b>	29th January 2024

**PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:**

Attach a completed copy of this sheet to each project (including multiple copies).	<input type="checkbox"/>
<b>Attach a Moodle submission receipt of the online project submission</b> , to each project (including multiple copies).	<input type="checkbox"/>
<b>You must ensure that you retain a HARD COPY of the project</b> , 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.

<b>Office Use Only</b>	
Signature:	
Date:	
Penalty Applied (if applicable):	

# Configuration Manual

Anant Sakharam Pednekar  
21188947

## 1 Introduction

This article is an in-depth guide for setting up a Kubernetes cluster on a Ubuntu server. It lays out the essential steps for building the cluster and deploying a robust Pod network in great detail and provides the essential technologies.

The manual's last section explores benchmarking tool configuration methods, shedding light on the complexities of k-bench and sysbench deployment. Insights gained from this investigation will help readers assess the efficiency and effectiveness of their Kubernetes environment, which can lead to better decisions on the allocation of resources.

## 2 Tools and Technologies Required

The following table outlines the essential tools and technologies utilized in this experiment, encompassing the clustering platform, application-container operating system, container software, container orchestrator, platform tools for monitoring, design specification, and other key elements necessary for the successful execution of the experiment.

Category	Details
Server	AWS EC2
Host Operating System (OS)	Ubuntu Server
Container Orchestrator	Kubernetes , Openshift
Tools for Benchmarking	Kbench, Benchmark-operator
Slave & Master Config	t4g.2xlarge (8 vCPUs and 32GB)

Table 1: Tools & Technologies

## 3 Creation of Kubernetes Cluster

This portion of the configuration manual utilizes AWS Cloud services, specifically EC2 instances, to construct a Kubernetes cluster. Refer 1.

The selected operating system for this research is Ubuntu Server. This approach is chosen for its compatibility and support in effortlessly creating Kubernetes clusters.

To set up the Kubernetes cluster, it is necessary to have both a master node and one or more slave nodes. For this benchmarking, the master and slave nodes are built using

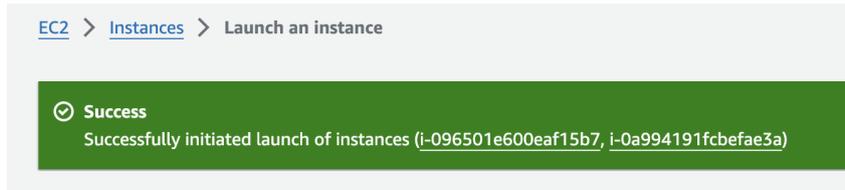


Figure 1: AWS Instances Launch

the t4g.2xlarge instance type, which has 8 virtual CPUs and 32GB of memory.

Before initiating the cluster, installing specific prerequisites on all nodes is necessary Refer 2 . Each node should perform the following commands <sup>1</sup>:

```
sudo apt-get install golang make tmux -y
export GOROOT=/usr/lib/go
export GOPATH=/home/ubuntu/go
```

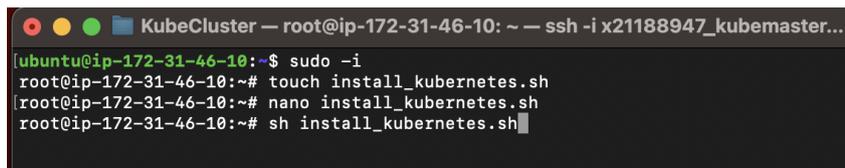


Figure 2: Install Kubernetes

The command in figure 3 is to be executed only on the master node :

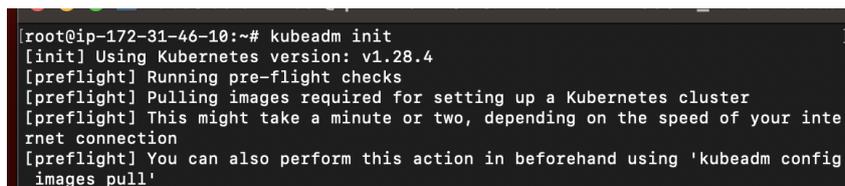


Figure 3: Kubernetes init

The master node will then create the cluster using kubeadm and generate a token. This token is essential for the slave/worker nodes to join the cluster. Refer 4

On each worker node, the following command should be executed using the token obtained from the master node :

---

<sup>1</sup>Link to install\_kubernetes.sh

```

KubeCluster — root@ip-172-31-46-10: ~ — ssh -i x21188947_kubemaster...

Your Kubernetes control-plane has initialized successfully!

To start using your cluster, you need to run the following as a regular user:

mkdir -p $HOME/.kube
sudo cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
sudo chown $(id -u):$(id -g) $HOME/.kube/config

Alternatively, if you are the root user, you can run:

export KUBECONFIG=/etc/kubernetes/admin.conf

You should now deploy a pod network to the cluster.
Run "kubectl apply -f [podnetwork].yaml" with one of the options listed at:
https://kubernetes.io/docs/concepts/cluster-administration/addons/

Then you can join any number of worker nodes by running the following on each as
root:

kubeadm join 172.31.46.10:6443 --token zss6gr.hx33usgkz9sdx27y \
--discovery-token-ca-cert-hash sha256:eadd8c1866fe88b46fd7e9d2083456f91c
e340c0df3765ffd9d7e4164416f28c
root@ip-172-31-46-10:~# █

```

Figure 4: Master Ready

```

kubeadm join 172.31.22.210:6443 \
--token mndegl.v4ez3pj9ru3i7qkp \
--discovery-token-ca-cert-hash \
sha256:a1fbbcb61c03305f43d352a8f988aebf4b1a1909c574d6ac313e9ce3509578ed

```

Once joined, the output should resemble the one shown in figure 5

```

KubeCluster — root@ip-172-31-36-99: ~ — ssh -i x21188947_kubemaste...

root@ip-172-31-36-99:~# kubeadm join 172.31.46.10:6443 --token zss6gr.hx33usgkz9
sdx27y \
[
--discovery-token-ca-cert-hash sha256:eadd8c1866fe88b46fd7e9d2083456f91c
e340c0df3765ffd9d7e4164416f28c
[preflight] Running pre-flight checks
[preflight] Reading configuration from the cluster...
[preflight] FYI: You can look at this config file with 'kubectl -n kube-system g
et cm kubeadm-config -o yaml'
[kubelet-start] Writing kubelet configuration to file "/var/lib/kubelet/config.y
aml"
[kubelet-start] Writing kubelet environment file with flags to file "/var/lib/ku
belet/kubeadm-flags.env"
[kubelet-start] Starting the kubelet
[kubelet-start] Waiting for the kubelet to perform the TLS Bootstrap...

This node has joined the cluster:
* Certificate signing request was sent to apiserver and a response was received.
* The Kubelet was informed of the new secure connection details.

Run 'kubectl get nodes' on the control-plane to see this node join the cluster.

root@ip-172-31-36-99:~# █

```

Figure 5: Worker init

To verify the functionality of cluster and install the Pod Network, the following commands are executed:

```

# Verify kubectl functionality
kubectl get nodes

```

```
# Install Pod Network (Weave)
kubectl apply -f \
https://github.com/weaveworks/weave
/releases/download/v2.8.1/weave-daemonset-k8s-1.9.yaml
```

Once done the output should resemble the one shown in figure 6

```
ubuntu@ip-172-31-46-10:~$ mkdir -p $HOME/.kube
sudo cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
sudo chown $(id -u):$(id -g) $HOME/.kube/config
ubuntu@ip-172-31-46-10:~$ kubectl get nodes
NAME                STATUS    ROLES    AGE   VERSION
ip-172-31-36-99     NotReady <none>   42s   v1.28.2
ip-172-31-46-10     NotReady control-plane 79s   v1.28.2
ubuntu@ip-172-31-46-10:~$ kubectl apply -f https://github.com/weaveworks/weave/releases/download/v2.8.1/weave-daemonset-k8s-1.9.yaml
serviceaccount/weave-net created
clusterrole.rbac.authorization.k8s.io/weave-net created
clusterrolebinding.rbac.authorization.k8s.io/weave-net created
role.rbac.authorization.k8s.io/weave-net created
rolebinding.rbac.authorization.k8s.io/weave-net created
daemonset.apps/weave-net created
ubuntu@ip-172-31-46-10:~$ kubectl get nodes
NAME                STATUS    ROLES    AGE   VERSION
ip-172-31-36-99     Ready     <none>   2m28s v1.28.2
ip-172-31-46-10     Ready     control-plane 3m5s   v1.28.2
ubuntu@ip-172-31-46-10:~$
```

Figure 6: Cluster Ready

## 4 Install Benchmarking operator

In this section, we focus on the installation process of the benchmarking operator, a crucial component for executing performance tests using sysbench within the Kubernetes environment.

The installation process is depicted in Figure 7.

```
ubuntu@ip-172-31-46-10:~$ cd benchmark-operator && make deploy
cd config/manager && /home/ubuntu/benchmark-operator/bin/kustomize edit set image controller=quay.io/cloud-bulldozer/benchmark-operator:latest
/home/ubuntu/benchmark-operator/bin/kustomize build config/default | kubectl apply -f -
namespace/benchmark-operator created
customresourcedefinition.apiextensions.k8s.io/benchmarks.ripsaw.cloudbulldozer.io created
serviceaccount/benchmark-operator created
clusterrolebinding.rbac.authorization.k8s.io/benchmark-operator-binding created
configmap/benchmark-manager-config created
deployment.apps/benchmark-controller-manager created
ubuntu@ip-172-31-46-10:~/benchmark-operator$
```

Figure 7: Install Benchmarking operator

### 4.1 Running sysbench tests

To initiate the sysbench tests, execute the following command. The YAML configuration file utilized is sourced from the standard templates provided by the benchmark operator.

```
kubectl apply -f config/samples/sysbench/cr.yaml
```

This command triggers the deployment of the sysbench workload with the specified configurations, allowing for comprehensive benchmarking within the Kubernetes cluster.

## 5 Install K-bench operator

This section details the installation process for the K-bench operator, an essential tool for conducting benchmark tests within a Kubernetes environment.

Before installing the K-bench operator, ensure that Kustomize is installed. Execute the following command to install Kustomize:

```
curl -s "https://raw.githubusercontent.com/kubernetes-sigs/\
kustomize/master/hack/install_kustomize.sh" | bash
```

Clone the K-bench repository from GitHub and navigate to the repository directory. Execute the installation script and move the K-bench executable to a directory included in the system's PATH.

```
git clone https://github.com/vmware-tanzu/k-bench.git
cd k-bench
go get -u github.com/kubernetes/kompose
go get golang.org/x/oauth2/google@v0.0.0-20200107190931-bf48bf16ab8d
./install.sh
```

```
sudo mv /home/ubuntu/go/bin/kbench
```

### 5.1 Running Tests with K-bench

To perform benchmark tests using K-bench, use the provided run scripts. Here are examples of running different tests:

```
./run.sh -r "kbench-run-dp_redis-1" -t "dp_redis" -o "./" \
| tee test-dp_redis.txt
./run.sh -r "kbench-run-dp_redis_service-1" -t "dp_redis_service" -o "./" \
| tee test-dp_redis_service.txt
./run.sh -r "kbench-run-cp_heavy_12client-1" -t "cp_heavy_12client" -o "./" \
| tee test-cp_heavy_12client.txt
```