

# Configuration Manual

MSc Research Project  
Cloud Computing

**Jialing Chen**  
Student ID: x23158131

School of Computing  
National College of Ireland

Supervisor: Shaguna Gupta

**National College of Ireland**  
**MSc Project Submission Sheet**



**School of Computing**

**Student Name:** ..... Jialing Chen  
**Student ID:** ..... x23158131  
**Programme:** ..... Cloud Computing **Year:** ..... 2024  
 ..... MSc Research Project  
**Module:** .....  
 ..... Shaguna Gupta  
**Lecturer:** .....  
**Submission Due Date:** ..... 01/09/2025  
**Project Title:** ..... Rule-based Task Scheduling in Cloud-Edge  
 ..... Computing for Energy and Resource Utilization Optimization  
 .....  
**Word Count:** .....2389.....

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:** ..... Jialing Chen  
 ..... 26<sup>th</sup> August, 2025  
**Date:** .....

**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

Jialing Chen  
23158131

## 1 Introduction

This manual explains how to set up and run the experiments for the rule-based task scheduling system in a cloud-edge environment. The setup uses basic tools like Minikube for creating a local Kubernetes cluster, and common monitoring tools like Prometheus and Grafana. Each step is explained in simple terms to make it easy to follow and repeat the experiments.

## 2 Environment Setup

### 2.1 Basic Requirements

Running these tests needs a good computer with enough power. The basic setup should have 8GB of memory and 4 processor cores to work well. While this seems like a lot of power, most modern computers have these features. If the computer isn't strong enough, the tests might be slow or have problems, since it needs to run both cloud and edge parts together.

The first program to install is Docker Desktop. This needs to be ready before anything else starts. Docker helps by putting programs into boxes that can run anywhere. It's like packing a suitcase with everything needed for a trip - once packed, it works the same way everywhere it goes.

Two more tools help run Kubernetes on the computer. Kubectl works like a remote control, sending commands to Kubernetes. Minikube creates a small test system that acts like a real cloud setup. This lets anyone try out different settings without spending money on big cloud services - just like having a small practice version at home.

### 2.2 Creating Kubernetes Cluster

Start by creating a test system with Minikube (Kubernetes.io, 2024) to make some computers act like cloud and edge machines:

1. Get Minikube running with the right amount of computer power:

```
minikube start --cpus=4 --memory=8192mb --nodes=2
```

2. Mark each computer as either cloud or edge:

```
kubectl label node minikube node-role="cloud"  
kubectl label node minikube-m02 edge-node="true"
```

3. Check that both nodes are ready:

```
kubectl get nodes --show-labels
```

```
minikube Ready control-plane 77d v1.33.1 beta.kubernetes.io/arch=arm64,beta.kubernetes.io/os=linux,edge-node=true,kubernetes.io/arch=arm64,kubernetes.io/hostname=minikube,kubernetes.io/os=linux,minikube.k8s.io/commit=f8f52f5de11fc6ad8244afac475e1d0f96841df1,minikube.k8s.io/name=minikube,minikube.k8s.io/primary=true,minikube.k8s.io/updated_at=2025_06_04T19_12_45_0700,minikube.k8s.io/version=v1.36.0,node-role.kubernetes.io/control-plane=node-role=cloud,node.kubernetes.io/exclude-from-external-load-balancers=
```

Figure 1: Node Labels Check Result

## 2.3 Setting Up Monitoring Tools

### 2.3.1 Installing Prometheus

Prometheus (Prometheus.io, 2024) helps collect important information about how the system is running:

1. Add Prometheus repository:

```
helm repo add prometheus-community \
https://prometheus-community.github.io/helm-charts
```

2. Create monitoring namespace:

```
kubectl create namespace monitoring
```

3. Install Prometheus:

```
helm install prometheus prometheus-community/kube-prometheus-stack \
-n monitoring
```

```
default prometheus-alertmanager-0 1/1 Running 113 2
5d
default prometheus-kube-state-metrics-57d654d7bf-5qzkt 1/1 Running 78 5
2d
default prometheus-prometheus-node-exporter-nn4kf 1/1 Running 107 (71m ago) 5
2d
default prometheus-prometheus-pushgateway-784c485d55-5dmgx 1/1 Running 32 5
2d
default prometheus-server-58984f5fd7-qdgnm 3/3 Running 75 (74m ago) 8
d
```

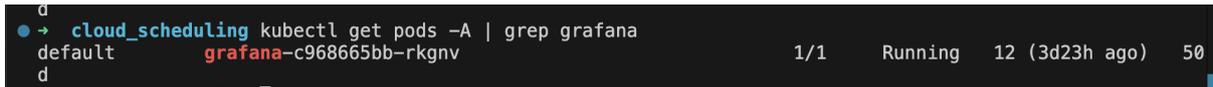
Figure 2: Prometheus Installation

### 2.3.2 Setting Up Grafana

Grafana (Grafana Labs, 2024) helps see the collected information in easy-to-understand graphs:

1. Access Grafana dashboard:

```
kubectl port-forward svc/prometheus-grafana 3000:80 -n monitoring
```



```
d
● → cloud_scheduling kubectl get pods -A | grep grafana
default grafana-c968665bb-rkgnv 1/1 Running 12 (3d23h ago) 50
d
```

Figure 3: Grafana Dashboard Access

Import basic dashboard templates for:

- CPU usage monitoring
- Memory usage tracking

## 3 Experiment Configurations

### 3.1 Experiment 1: Task Migration (R1)

This experiment shows how tasks move between cloud and edge nodes based on CPU usage.

#### 3.1.1 Core Application Setup

The main program is built using Flask and runs different types of tasks. It has simple commands to start and stop heavy workloads. The program is packed into a container to make it easy to run anywhere in the system.

```

1  from flask import Flask
2  import subprocess
3  import os
4  import time
5
6  app = Flask(__name__)
7
8  @app.route('/')
9  def index():
10     return "Welcome to the Cloud Node Service"
11
12 @app.route('/compute-heavy')
13 def compute_heavy():
14     try:
15         num_cores = os.cpu_count()
16         # print(f"Number of cores: {num_cores}")
17         # target_cores =int(num_cores)
18         duration = 600
19         subprocess.Popen(['stress-ng', '--cpu', str(num_cores), '--cpu-method', 'matrixprod', '--cpu-
20             return f"Started stress with {num_cores} threads for {duration} seconds!", 200
21
22     except Exception as e:
23         return f"Error occurred: {str(e)}", 500
24
25 @app.route('/stop')
26 def stop_stress():
27     global stress_process
28     if stress_process and stress_process.poll() is None:
29         stress_process.terminate()
30         stress_process.wait(timeout=5)
31         stress_process = None

```

Figure 4: Core Application Code

### 3.1.2 Cloud Node Setup

The cloud deployment configuration specifies resource limits and node selection rules. The deployment-cloud.yaml sets CPU limits to 8 cores and minimum requests to 1 core, ensuring proper resource allocation for the workload tests.

```

1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: flask-task-cloud
5    labels:
6      app: flask-task
7  spec:
8    replicas: 1
9    selector:
10     matchLabels:
11       app: flask-task
12   template:
13     metadata:
14       labels:
15         app: flask-task
16     spec:
17       nodeSelector:
18         node-role: cloud
19       containers:
20         - name: flask-container
21           image: flask-tasks-service
22           imagePullPolicy: IfNotPresent
23           ports:
24             - containerPort: 5000
25           resources:
26             requests:
27               cpu: 1
28             limits:
29               cpu: 8
30
31

```

Figure 5: Cloud Node Deployment Configuration

### 3.1.3 Monitoring Configuration

Prometheus handles the monitoring setup with custom alert rules. The `cpu-alert-rules.yaml` defines when migration should trigger based on CPU usage. The `prometheus-config.yaml` sets up the data collection rules.

```

1  groups:
2    - name: cloud-cpu-alert
3      rules:
4        - alert: HighCloudCPUUsage
5          expr: (1 - avg(rate(node_cpu_seconds_total{mode="idle"}[2m])) by (instance)) * 100 >= 50
6          for: 10s
7          labels:
8            severity: critical
9          annotations:
10             summary: "Cloud node CPU >= 50%"
11             description: "CPU usage has exceeded 50% on {{ $labels.instance }} for 10 seconds"
12

```

Figure 6: CPU Alert Rules Configuration

### 3.1.4 Migration Control

The webhook system manages task migration between nodes. The `webhook-deploy.yaml` and `webhook-rbac.yaml` set up the migration service with proper permissions. A custom webhook handler processes migration requests when CPU thresholds are exceeded.

```
1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: webhook-handler
5  spec:
6    replicas: 1
7    selector:
8      matchLabels:
9        app: webhook
10   template:
11     metadata:
12       labels:
13         app: webhook
14     spec:
15       serviceAccount: webhook-sa
16       containers:
17         - name: webhook
18           image: webhook-handler
19           imagePullPolicy: IfNotPresent
20           ports:
21             - containerPort: 5001
22
23
24 ---
25  apiVersion: v1
26  kind: Service
27  metadata:
28    name: webhook-handler-service
29  spec:
30    selector:
31      app: webhook
```

Figure 7: Webhook Deployment Configuration

### 3.1.5 Service Configuration

The `service-cloud.yaml` defines how the application is exposed within the cluster. It sets up the necessary port mappings and service discovery rules for the migration system to work properly.

```
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: flask-task-service
5  spec:
6    selector:
7      app: flask-task
8    ports:
9      - protocol: TCP
10       port: 5000
11       targetPort: 5000
12  %L to chat, %K to generate
```

Figure 8: Service Configuration

## 3.2 Experiment 2: Load Balancing (R2)

This experiment uses Istio (Istio.io, 2024), a powerful service mesh platform, to direct traffic based on user location. Istio helps manage network traffic and implement location-based routing rules in a cloud-edge environment.

### 3.2.1 Basic Service Setup

Programs run on both cloud and edge computers. Each one gets different power limits. Cloud computers can do more heavy work with 8 processors, while edge computers handle smaller tasks. This helps share work across different places effectively.

```
1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: flask-task-cloud-lb
5    labels:
6      app: flask-cloud-lb
7  spec:
8    replicas: 1
9    selector:
10     matchLabels:
11       app: flask-task-lb
12       version: cloud
13   template:
14     metadata:
15       labels:
16         app: flask-task-lb
17         version: cloud
18     spec:
19       nodeSelector:
20         node-role: cloud
21       containers:
22         - name: flask-container-cloud-lb
23           image: flask-tasks-service:latest
24           imagePullPolicy: IfNotPresent
25         ports:
26           - containerPort: 5000
27       resources:
28         requests:
29           cpu: 1
```

Figure 9: Cloud Server Configuration

```

1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: flask-task-edge-lb
5    labels:
6      app: flask-edge-lb
7  spec:
8    replicas: 1
9    selector:
10   matchLabels:
11     app: flask-task-lb
12     version: edge
13   template:
14     metadata:
15       labels:
16         app: flask-task-lb
17         version: edge
18     spec:
19       nodeSelector:
20         edge-node: "true"
21       containers:
22         - name: flask-container-edge-lb
23           image: flask-tasks-service:latest
24           imagePullPolicy: IfNotPresent
25           ports:
26             - containerPort: 5000
27           resources:
28             requests:|
29             cpu: 1

```

Figure 10: Edge Server Configuration

### 3.2.2 Traffic Routing Setup

Basic rules help control where requests go in the network. Each request gets checked to see where it came from. US requests go to nearby edge computers for faster service. Everything else gets sent to cloud computers for processing.

```

1  apiVersion: networking.istio.io/v1beta1
2  kind: VirtualService
3  metadata:
4    name: flask-task-vs
5  spec:
6    hosts:
7    - "*"
8    gateways:
9    - flask-gateway
10   http:
11   - match:
12     - headers:
13       x-user-location:
14         exact: "US"
15     route:
16     - destination:
17       host: flask-task-service-lb
18       subset: edge
19       port:
20         number: 80
21   - route:
22     - destination:
23       host: flask-task-service-lb
24       subset: cloud
25       port:
26         number: 80
27
28
29

```

Figure 11: Traffic Routing Rules

### 3.2.3 Server Group Settings

Computers get sorted into cloud or edge groups. This grouping helps control how work moves between different parts of the network. The settings make sure requests go to the right group of servers and keep everything running smoothly.

```

1  apiVersion: networking.istio.io/v1beta1
2  kind: DestinationRule
3  metadata:
4    name: flask-task-dr
5  spec:
6    host: flask-task-service-lb
7    subsets:
8    - name: cloud
9      labels:
10     version: cloud
11    - name: edge
12      labels:
13     version: edge
14    trafficPolicy:
15      tls:
16        mode: DISABLE
17    %L to chat, %K to generate
18
19

```

Figure 12: Server Group Configuration

### 3.2.4 Network Gateway Setup

The gateway works like a front door for all incoming traffic. It helps manage how requests come into the system and makes sure they get sent to the right place. This setup helps balance the work across all servers.

```
1  apiVersion: networking.istio.io/v1beta1
2  kind: Gateway
3  metadata:
4    name: flask-gateway
5  spec:
6    selector:
7      istio: ingressgateway
8    servers:
9      - port:
10         number: 80
11         name: http
12         protocol: HTTP
13         hosts:
14           - "*"

```

Figure 13: Gateway Configuration

### 3.2.5 Service Connection Setup

The service settings tell different parts of the system how to talk to each other. They set up the right connections and make sure everything can work together properly. This helps keep the whole system running smoothly.

```
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: flask-task-service-lb
5  spec:
6    selector:
7      app: flask-task-lb
8    ports:
9      - port: 80
10        targetPort: 5000
11        protocol: TCP
12        name: http
13
14
15  | %! to chat, %! to generate

```

Figure 14: Service Connection Settings

## 3.3 Experiment 3: Cache Policy (R3)

This test uses Istio (Istio.io, 2024) to handle product requests in a smarter way. When someone looks for a popular product, Istio helps send their request to the right place -

usually to an edge server that already has the product information saved. This makes things faster because the information is closer to where it's needed.

### 3.3.1 Basic Product Service Setup

The system uses a simple product service that shows different items. It runs on port 5050 and can tell if a product is popular (hot) or normal. The service keeps track of which products are accessed more often.

```
1  from flask import Flask, request
2  import time
3
4  app = Flask(__name__)
5
6  HOT_PRODUCTS = {"101"}
7
8  @app.route("/product/<product_id>")
9  def product_page(product_id):
10     client_ip = request.remote_addr
11     is_hot = product_id in HOT_PRODUCTS
12
13     print(f"[Access] Product ID: {product_id} | Hot: {is_hot} | From: {client_ip}")
14
15     return f"""
16     <html>
17     <h2>Product Page: {product_id}</h2>
18     <p>Status: {"Hot (Edge Cached)" if is_hot else "Normal (Cloud Cached)"}</p>
19     </html>
20     """
21
22 if __name__ == "__main__":
23     app.run(host="0.0.0.0", port=5050)
24     %%L to chat, %%K to generate
25
26
27
28
```

Figure 15: Product Service Implementation

### 3.3.2 Cache Location Rules

The system has special rules for storing product information. Popular products are saved closer to users (on edge servers), while other products stay in the main cloud storage. This helps get popular items to users faster.

```

1  apiVersion: networking.istio.io/v1beta1
2  kind: VirtualService
3  metadata:
4    name: product-service-vs
5  spec:
6    hosts:
7    - "*"
8    gateways:
9    - product-gateway
10   http:
11   - match:
12     - uri:
13       prefix: "/product/101"
14     route:
15     - destination:
16       host: product-service-cp
17       subset: edge
18       port:
19         number: 80
20   - match:
21     - uri:
22       prefix: "/product/"
23     route:
24     - destination:
25       host: product-service-cp
26       subset: cloud
27       port:
28         number: 80
29

```

Figure 16: Cache Location Rules Configuration

### 3.3.3 Server Deployment Setup

The system runs on both cloud and edge servers. Each server has its own job - edge servers handle popular items, while cloud servers handle everything else. This split helps manage product information better.

```
1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: product-cloud-cp
5  spec:
6    replicas: 1
7    selector:
8      matchLabels:
9        app: product-task-cp
10       version: cloud
11   template:
12     metadata:
13       labels:
14         app: product-task-cp
15         version: cloud
16     spec:
17       nodeSelector:
18         node-role: cloud
19       containers:
20         - name: product-container-cloud
21           image: product-service:latest
22           imagePullPolicy: IfNotPresent
23           ports:
24             - containerPort: 5050
25           resources:
26             requests:
27               cpu: 1
28             limits:
29               cpu: 8
```

Figure 17: Cloud Server Deployment Configuration

```

1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: product-edge-cp
5  spec:
6    replicas: 1
7    selector:
8      matchLabels:
9        app: product-task-cp
10       version: edge
11   template:
12     metadata:
13       labels:
14         app: product-task-cp
15         version: edge
16     spec:
17       nodeSelector:
18         edge-node: "true"
19       containers:
20         - name: product-container-edge
21           image: product-service:latest
22           imagePullPolicy: IfNotPresent
23           ports:
24             - containerPort: 5050
25           resources:
26             requests:
27               cpu: 1
28             limits:
29               cpu: 8

```

Figure 18: Edge Server Deployment Configuration

### 3.3.4 Traffic Management

The system uses a gateway to handle incoming requests. When someone looks for a product, the gateway checks if it's a popular item (like product 101) and sends the request to the right server. This helps get product information faster.

```

1  apiVersion: networking.istio.io/v1beta1
2  kind: Gateway
3  metadata:
4    name: product-gateway
5  spec:
6    selector:
7      istio: ingressgateway
8    servers:
9      - port:
10         number: 80
11         name: http
12         protocol: HTTP
13       hosts:
14         - "*"

```

Figure 19: Gateway Traffic Management Configuration

### 3.3.5 Connection Settings

The service settings make sure all parts of the system can talk to each other. They help connect the cloud and edge servers and make sure product information flows smoothly between them.

```
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: product-service-cp
5  spec:
6    selector:
7      app: product-task-cp
8    ports:
9      - port: 80
10     targetPort: 5050
11     protocol: TCP
12     name: http
13  ☐ %L to chat, %K to generate
```

Figure 20: Service Connection Configuration

```
1  apiVersion: networking.istio.io/v1beta1
2  kind: DestinationRule
3  metadata:
4    name: product-service-dr
5  spec:
6    host: product-service-cp
7    subsets:
8      - name: cloud
9        labels:
10         version: cloud
11      - name: edge
12        labels:
13         version: edge
14    trafficPolicy:
15      tls:
16        mode: DISABLE
17
18
19  ☐ %L to chat, %K to generate
```

Figure 21: Destination Rules Configuration

## 3.4 Experiment 4: Cloud Scheduling (R4)

This test uses HPA (Horizontal Pod Autoscaler) (Kubernetes.io/HPA, 2024) to add or remove pods automatically when needed. HPA watches how busy the system is and helps make sure there are enough pods to handle all the work, kind of like opening more checkout counters in a store when it gets crowded.

### 3.4.1 Basic Service Setup

The system runs a service that can create different levels of work. It has simple commands to start and stop tasks, and can measure how busy the computer is. The service runs on port 6060 and can handle multiple tasks at once.

```
1  from flask import Flask
2  import multiprocessing
3  import time
4  import os
5
6  app = Flask(__name__)
7  | TAB to jump here
8  processes = []
9  load_started = False
10
11 def cpu_burner():
12     while True:
13         start = time.time()
14         while time.time() - start < 0.08:
15             pass
16             time.sleep(0.02)
17
18 @app.route("/start")
19 def start_load():
20     global load_started, processes
21     %L to chat, %K to generate
22     if load_started:
23         return "Load already running.\n"
24
25     num_procs = 2
26     load_started = True
27
28     for _ in range(num_procs):
29         p = multiprocessing.Process(target=cpu_burner)
30         p.start()
```

Figure 22: Service Implementation

### 3.4.2 Automatic Scaling Rules

The system can grow or shrink based on how busy it gets. When the computer gets too busy (over 20% CPU use), it can add more space to handle the work. It can use between 1 and 2 copies of the service to handle different amounts of work.

```

1  apiVersion: autoscaling/v2
2  kind: HorizontalPodAutoscaler
3  metadata:
4    name: scale-hpa
5  spec:
6    scaleTargetRef:
7      apiVersion: apps/v1
8      kind: Deployment
9      name: scale-task-cloud
10   minReplicas: 1
11   maxReplicas: 2
12   metrics:
13     - type: Resource
14       resource:
15         name: cpu
16         target:
17           type: Utilization
18           averageUtilization: 20
19
20

```

Figure 23: HPA Scaling Rules Configuration

### 3.4.3 Work Distribution Setup

The system splits work between different servers. It uses special rules to decide which server should handle each task. This helps make sure no single server gets too busy.

```

1  apiVersion: networking.istio.io/v1beta1
2  kind: VirtualService
3  metadata:
4    name: scale-service-vs
5  spec:
6    hosts:
7      - "*"
8    gateways:
9      - scale-gateway
10   http:
11     - match:
12       - headers:
13         x-user-location:
14           exact: "edge"
15       route:
16         - destination:
17             host: scale-service
18             subset: edge
19             port:
20               number: 8080
21         - route:
22             - destination:
23                 host: scale-service
24                 subset: cloud
25             port:
26               number: 80
27
28
29

```

Figure 24: Work Distribution Configuration

### 3.4.4 Server Configuration

The system runs on both cloud and edge servers. Each server has its own settings for how much work it can handle. The cloud server can handle bigger tasks, while the edge server takes care of smaller ones.

```
1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: scale-task-cloud
5  spec:
6    replicas: 1
7    selector:
8      matchLabels:
9        app: scale-task
10       version: cloud
11   template:
12     metadata:
13       labels:
14         app: scale-task
15         version: cloud
16     spec:
17       nodeSelector:
18         node-role: cloud
19       containers:
20         - name: scale-container-cloud
21           image: scale-service:latest
22           imagePullPolicy: IfNotPresent
23           ports:
24             - containerPort: 6060
25           resources:
26             requests:
27               cpu: "100m"
28             limits:
29               cpu: 5
```

Figure 25: Cloud Server Configuration

```

1  apiVersion: apps/v1
2  kind: Deployment
3  metadata:
4    name: scale-task-edge
5  spec:
6    replicas: 1
7    selector:
8      matchLabels:
9        app: scale-task
10       version: edge
11   template:
12     metadata:
13       labels:
14         app: scale-task
15         version: edge
16     spec:
17       nodeSelector:
18         edge-node: "true"
19       containers:
20         - name: scale-container-edge
21           image: scale-service:latest
22           imagePullPolicy: IfNotPresent
23           ports:
24             - containerPort: 6060
25           resources:
26             requests:
27               cpu: "100m"
28             limits:
29               cpu: 5

```

Figure 26: Edge Server Configuration

### 3.4.5 Connection Management

Basic rules help all parts of the network talk to each other. This makes work flow easily between computers and keeps everything running smoothly.

```

1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: scale-service
5  spec:
6    selector:
7      app: scale-task
8    ports:
9      - port: 80
10       targetPort: 6060
11       protocol: TCP
12       name: http
13  ⌘L to chat, ⌘K to generate

```

Figure 27: Service Connection Configuration

```

1  apiVersion: networking.istio.io/v1beta1
2  kind: DestinationRule
3  metadata:
4    name: scale-service-dr
5  spec:
6    host: scale-service
7    subsets:
8      - name: cloud
9        labels:
10         version: cloud
11      - name: edge
12        labels:
13         version: edge
14    trafficPolicy:
15      tls:
16        mode: DISABLE
17

```

Figure 28: Destination Rules Configuration

## 4 Running the Experiments

### 4.1 Task Migration Test (R1)

The first test checks how work moves between cloud and edge when computers get busy. Here are the simple steps:

#### 4.1.1 Getting ready

Run programs on cloud computers

```
kubectl apply -f task-migration/deployment-cloud.yaml
```

```
kubectl apply -f task-migration/service-cloud.yaml
```

Check if everything starts correctly

```
kubectl get pods -o wide
```

#### 4.1.2 Starting the test

- Visit <http://localhost:5000/start-load> in a browser
- Give it half a minute to start working
- Watch the Grafana screen as processor use increases

#### 4.1.3 Looking at results

- Work moves to different computers when processor use reaches 50
- Grafana shows how busy each computer gets
- Check `kubectl get pods -o wide` to find where programs are running

## 4.2 Load Balancing Test (R2)

The second test shows how requests from different locations get handled. Follow these steps:

### 4.2.1 First steps

Run test programs on cloud and edge computers

```
kubectl apply -f load-balancing/deployment-cloud-lb.yaml
```

```
kubectl apply -f load-balancing/deployment-edge-lb.yaml
```

```
kubectl apply -f load-balancing/service-lb.yaml
```

### 4.2.2 Check different places

Send test requests from US locations

```
curl -H "x-user-location: US" http://localhost:8080/start-load
```

Send test requests from other places

```
curl http://localhost:8080/start-load
```

### 4.2.3 Check what happens

- Requests from the US go to the edge node
- Other requests go to the cloud node
- Grafana shows how requests get split up

## 4.3 Cache Policy Test (R3)

This test checks how the system handles popular products. Here's the process:

### 4.3.1 Setting up

```
kubectl apply -f cache-policy/deployment-cloud-cp.yaml
```

```
kubectl apply -f cache-policy/deployment-edge-cp.yaml
```

```
kubectl apply -f cache-policy/service-cp.yaml
```

### 4.3.2 Testing different products

Try getting a popular product (product 101)

```
http://localhost:5050/product/101
```

Try getting a regular product (product 102)

```
http://localhost:5050/product/102
```

### 4.3.3 Seeing what happens

- Popular product (101) requests go to the edge node
- Regular product requests go to the cloud node
- Grafana shows the speed difference between both

## 4.4 Cloud Scheduling Test (R4)

This test shows how the system adjusts resources when busy. Here's how to do it:

### 4.4.1 Getting it running

```
kubect1 apply -f cloud-scheduling/deployment-cloud-cs.yaml
kubect1 apply -f cloud-scheduling/service-cpu-load.yaml
kubect1 apply -f cloud-scheduling/hpa-scale.yaml
```

### 4.4.2 Making the system busy

- Go to <http://localhost:6060/start> to get things going
- Wait a few minutes to let the load build up
- Check how many pods are running with:

```
kubect1 get hpa
kubect1 get pods
```

### 4.4.3 Watching what happens

- When CPU hits 20%, the system adds more pods
- Grafana shows new pods being created
- When things quiet down, extra pods go away

## 4.5 Important Tips

- Take a break between tests to let the system settle down
- Clean up old test stuff before starting new tests
- Keep an eye on Grafana charts to spot any problems
- If something goes wrong, check the pod logs:

```
kubect1 logs <pod-name>
```

## 5 Troubleshooting Tips

Here are some common problems that might pop up during the experiments and how to fix them:

### 5.1 Pod Problems

If pods won't start:

```
kubect1 describe pod <pod-name>
kubect1 get nodes
```

Look for things like not enough CPU or memory on the nodes. Sometimes just waiting a minute and trying again helps.

## 5.2 Monitoring Issues

If Grafana isn't showing any data:

```
kubectl get pods -n monitoring
kubectl logs -n monitoring prometheus-grafana-<pod-id>
```

Make sure all monitoring pods are running. Sometimes restarting Prometheus helps:

```
kubectl rollout restart deployment prometheus-server -n monitoring
```

## 5.3 Traffic Routing Problems

If requests aren't going to the right place:

```
kubectl get virtualservice
kubectl describe virtualservice
```

Check if the Istio rules are set up right. The x-user-location header needs to be exactly "US" (uppercase) for edge routing.

## 5.4 Cache Problems

If product 101 isn't faster than others:

- Check if edge pods are running
- Look at the routing rules
- Try restarting the edge service:

```
kubectl rollout restart deployment edge-deployment
```

## 5.5 Scaling Issues

If pods aren't scaling automatically:

```
kubectl describe hpa
kubectl top pods
```

Check if the metrics server is working and CPU usage is being measured correctly.

## 5.6 Quick Fixes

Most problems can be fixed by:

- Waiting a minute or two
- Deleting pods and letting them restart
- Checking the logs with `kubectl logs`
- Making sure all services are running
- Restarting the specific service that's having trouble

If these steps don't help, try cleaning up everything and starting fresh:

```
kubectl delete -f <problem-deployment-file>
kubectl apply -f <problem-deployment-file>
```

## 6 References

Kubernetes.io. (2024). Minikube - Kubernetes Local Machine Solution. [online] Available at: <https://minikube.sigs.k8s.io/docs/> [Accessed 15 Mar. 2024].

Prometheus.io. (2024). Prometheus - Monitoring system & time series database. [online] Available at: <https://prometheus.io/docs/introduction/overview/> [Accessed 15 Mar. 2024].

Grafana Labs. (2024). Grafana - The open observability platform. [online] Available at: <https://grafana.com/docs/grafana/latest/> [Accessed 15 Mar. 2024].

Istio.io. (2024). Istio / What is Istio?. [online] Available at: <https://istio.io/latest/about/service-mesh/> [Accessed 15 Mar. 2024].

Kubernetes.io/HPA. (2024). Horizontal Pod Autoscaling. [online] Available at: <https://kubernetes.io/docs/tasks/run-application/horizontal-pod-autoscale/> [Accessed 15 Mar. 2024].