Konténer orkesztráció és autóskálázás MiCADO referencia architektúrával

Container orchestration and autoscaling by MiCADO

Kovács József

HUN REN | SZTAKI

Motivation

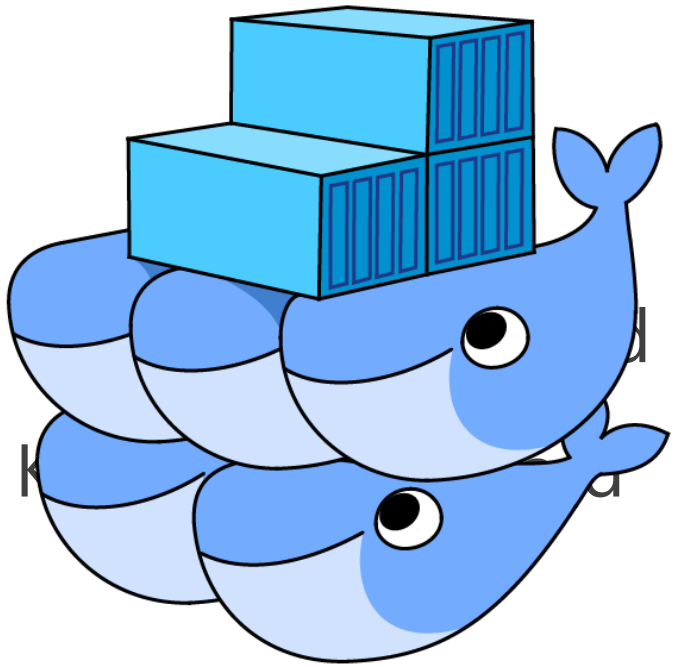

Kubernetes is good

for Cloud Native Microservices Architectures

- Self-healing
- Auto-scaling
- Health-checking
- Rolling updates
- Networking
- Security

Kubernetes is hard

(even when using a managed service)

- Deploying/managing a cluster
- Configuring a cluster
- Understanding abstractions
  - pod, job, deployment, replica set
- Writing templates
  - manifest files
- Debugging

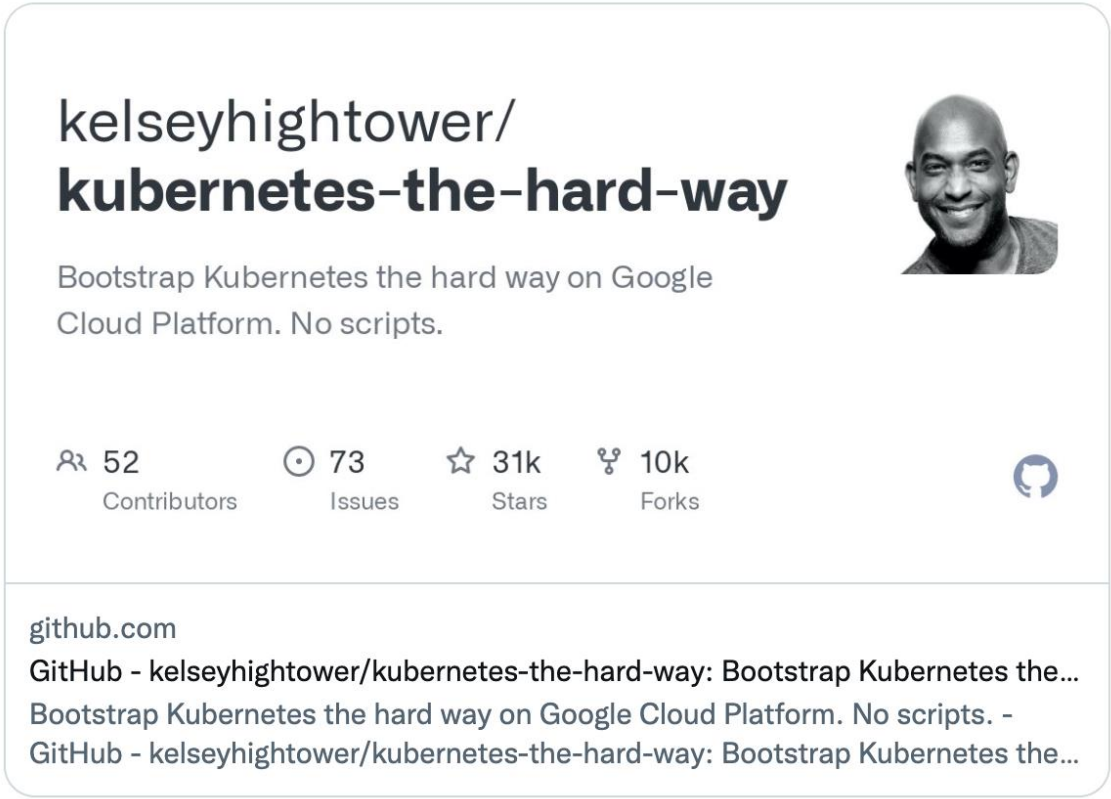

Application level orchestration

- multiple heterogeneous clouds
- wide range of scaling policies
- wide range of monitoring parameters
- advanced security solutions

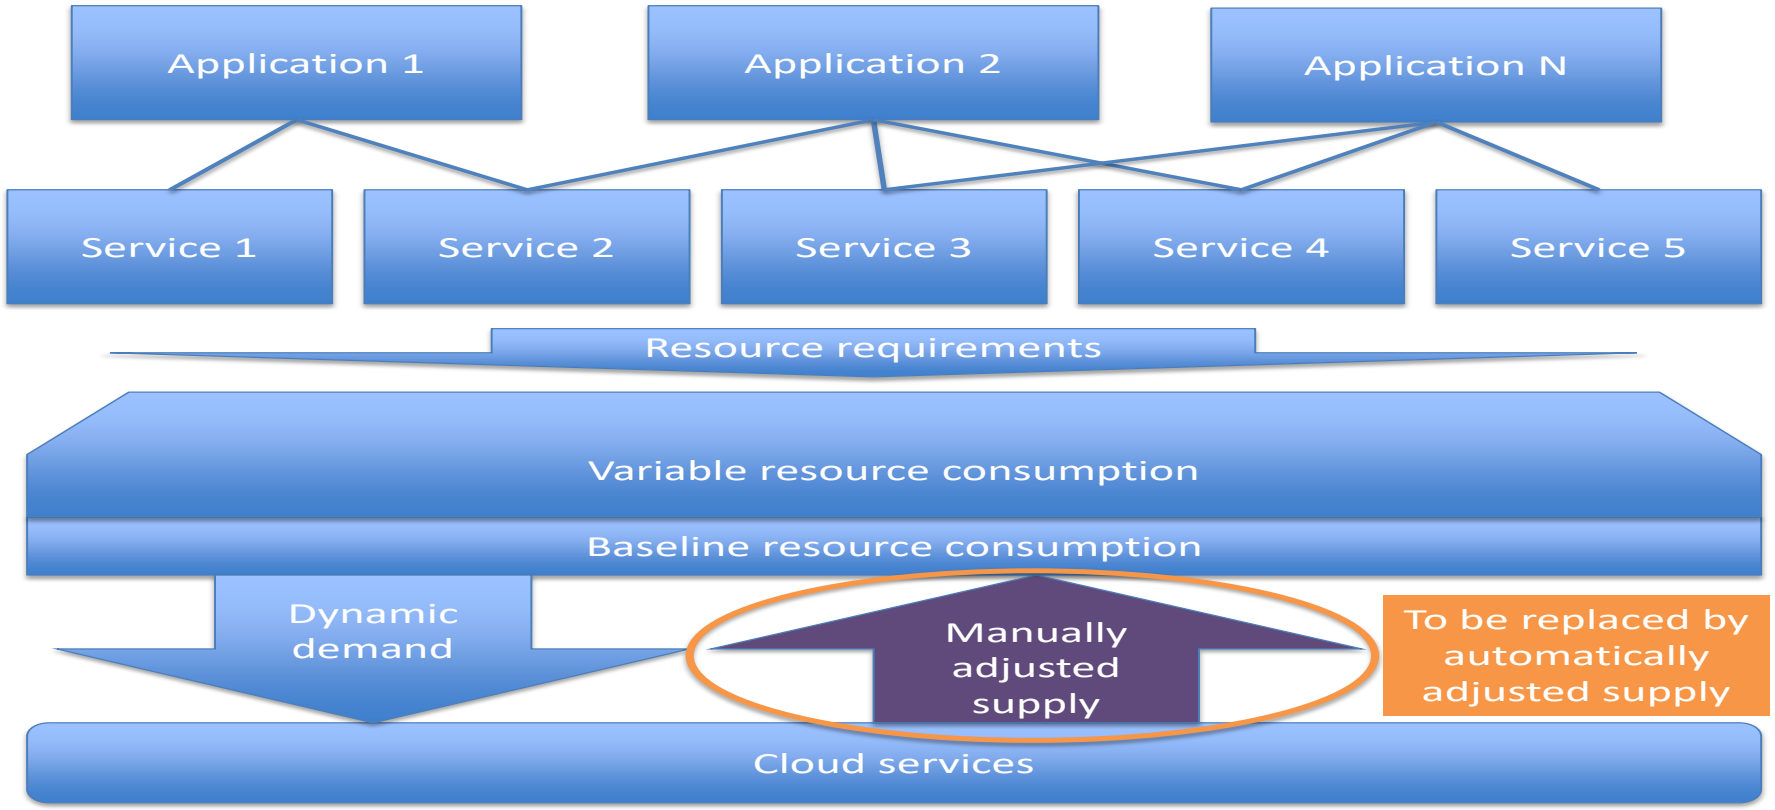

To achieve resource scalability and efficient resource utilisation supporting

Solution

Dynamic Cloud Orchestrator

- "One-click" deployment of an enhanced Kubernetes cluster
- Deploys, provisions, manages (auto-scaling, self-healing):
  - Applications (containers)
  - Cloud resources (virtual machines)
- Improved security
- Metrics dashboard

MiCADO – Microservices-based Cloud Application-level Dynamic Orchestrator

- History
  - Result of the **H2020 COLA** (Cloud Orchestration at the Level of Application) project (2017-2019)
  - Based on cooperation between Westminster University and SZTAKI
  - Since 2019, Westminster University has taken over the maintenance and development, SZTAKI contribution
  - Further developed in many European projects since the first prototype
    - currently actively developed in the H2020 DIGITbrain project
    - used in PITHIA-NRF, CO-VERSATILE (and Harpocrates and ARCAFF from October 2022)

MiCADO – Microservices-based Cloud Application-level Dynamic Orchestrator

- Main features
  - Automated application **deployment** based on TOSCA-based application description templates
  - **TOSCA -** Topology and Orchestration Specification for **Cloud** Applications
  - Automated **scaling** based on highly customisable scaling policies
    - scaling at both **container and virtual machine** levels
  - **Multi-cloud** support – application portability
  - Policy driven **security** settings
  - **Open source** and fully managed distributions
  - **Job queue** management with the extension of JQueuer
  - **Edge support** introduced for IoT-based applications

The overall concept

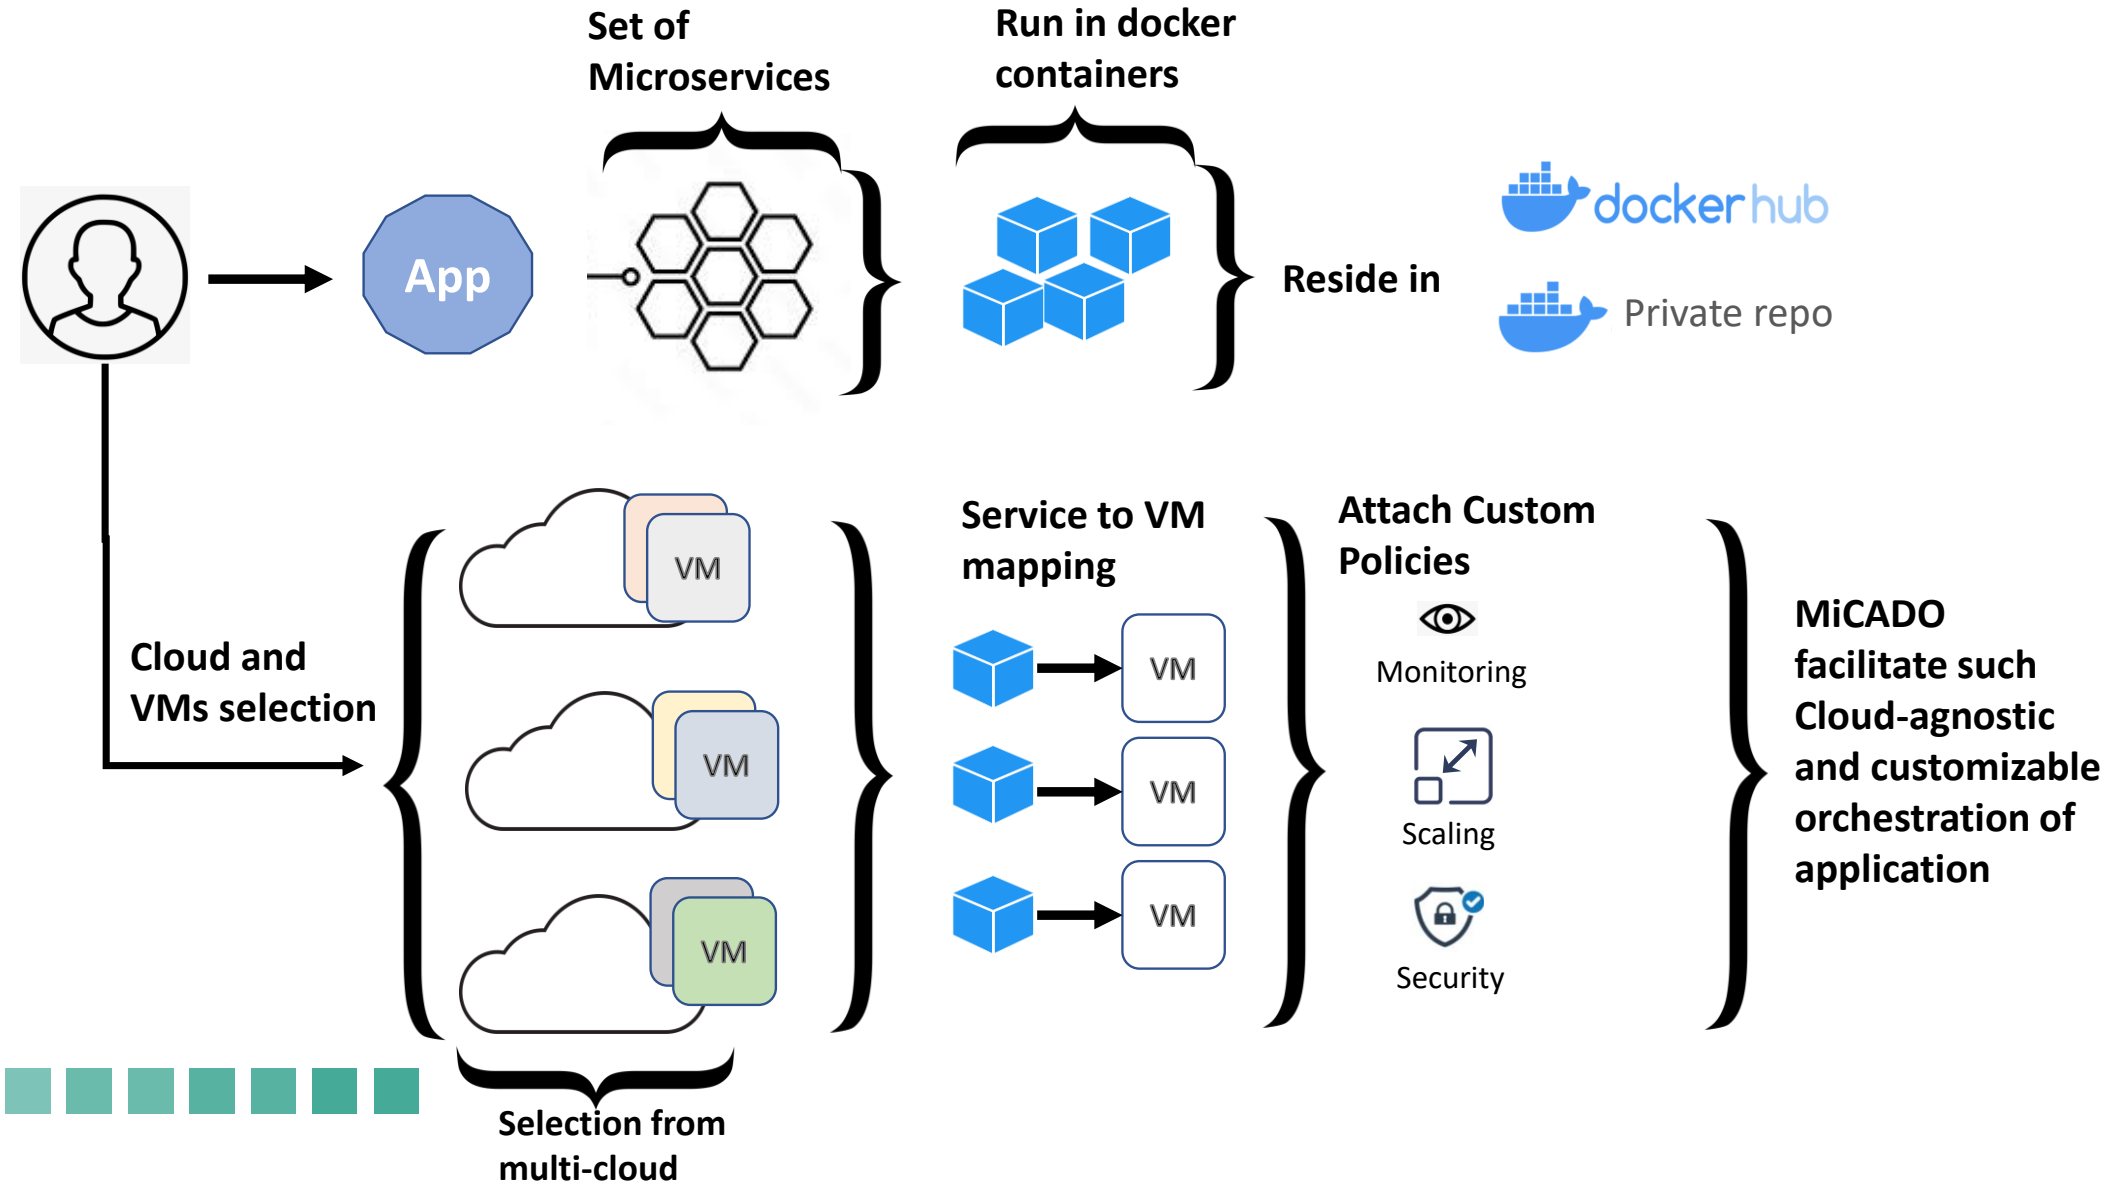

High-level architecture

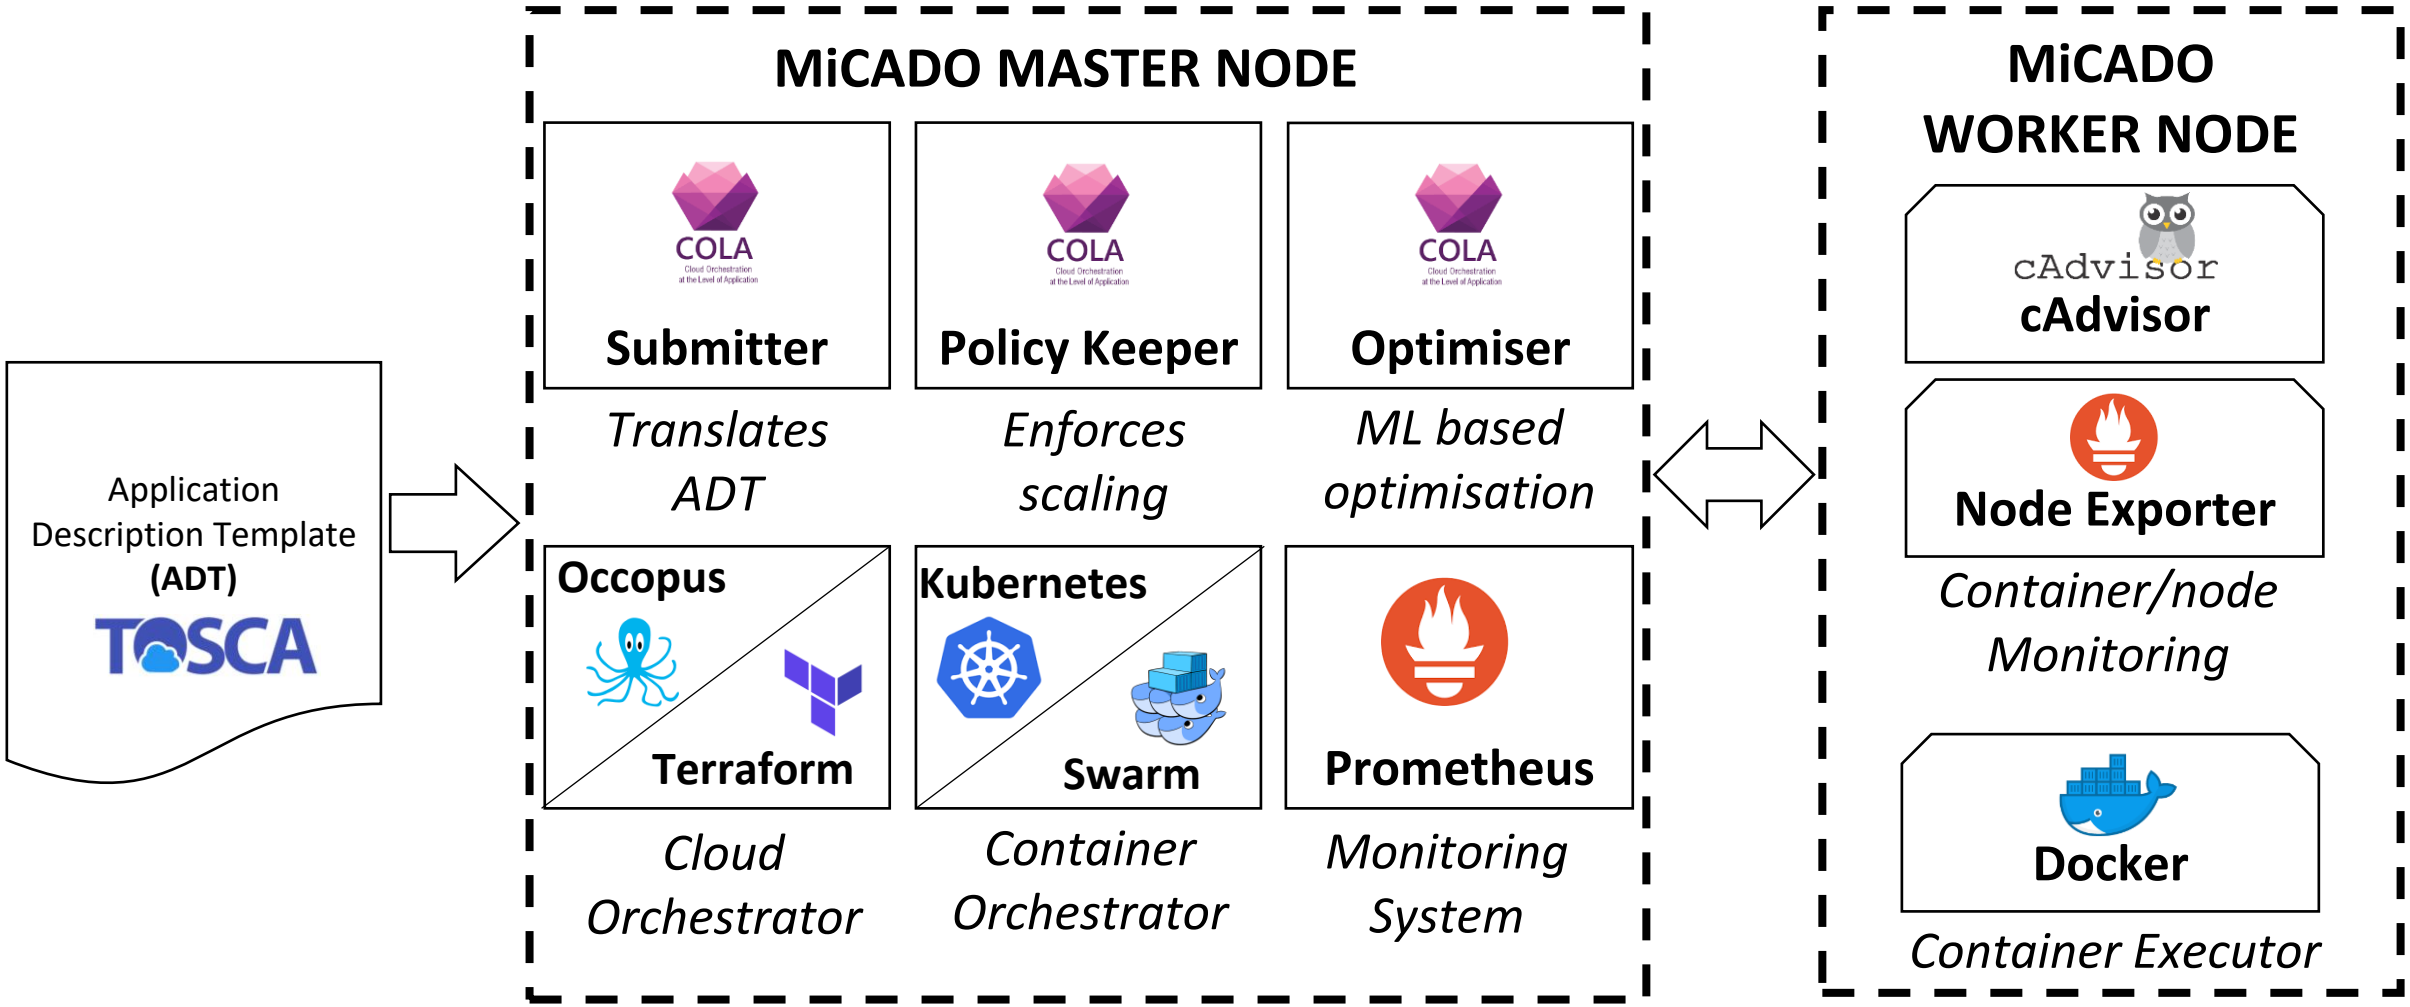

Application Description - Basics

Cloud Infrastructure (Instance size, SSH keys, opened ports, VM image)

```
YOUR-VIRTUAL-MACHINE:
  type: tosca.nodes.MiCADO.Nova.Compute
  properties:
    image_id: ADD_YOUR_ID_HERE (e.g. d4f4e496-031a-4f49-b034-f8dafe28e01c)
    flavor_name: ADD_YOUR_ID_HERE (e.g. 3)
    project_id: ADD_YOUR_ID_HERE (e.g. a678d20e71cb4b9f812a31e5f3eb63b0)
    network_id: ADD_YOUR_ID_HERE (e.g. 3fd4c62d-5fbe-4bd9-9a9f-c161dabeefde)
    key_name: ADD_YOUR_KEY_HERE (e.g. keyname)
    security_groups:
      - ADD_YOUR_ID_HERE (e.g. d509348f-21f1-4723-9475-0cf749e05c33)

  interfaces:
    Occopus:
      create:
        inputs:
          endpoint: ADD_YOUR_ENDPOINT (e.g https://sztaki.cloud.mta.hu:5000/v3)
```

Container Infrastructure (Container, volumes, configurations)

```
vm-node:
  type: tosca.nodes.MiCADO.EC2.Compute
  ...(truncated)...

app-container:
  type: tosca.nodes.MiCADO.Container.Application.Docker
  properties:
    image: nginx
  requirements:
  - host: vm-node
  interfaces:
    Kubernetes:
      create:
```

Application Description - Policies

- Monitoring subsystem
  - Monitoring metrics are collected by dynamically attachable data collectors (Prometheus exporters)
  - System and application metrics
- Highly customisable scaling system
  - Scaling of BOTH containers and virtual machines are supported
  - Scaling logic is fully programmable (using Python)
  - Various strategies (load-based, deadline-based, event-based, Scheduled)

```
policies:
  - monitoring:
      type: tosca.policies.Monitoring.MiCADO
      properties:
        enable_container_metrics: true
        enable_node_metrics: true
  - scalability:
      type: tosca.policies.Scaling.MiCADO.Container.CPU
      targets: [ stressng ]
      properties:
        constants:
          SERVICE_NAME: 'stressng'
          SERVICE_TH_MAX: '60'
          SERVICE_TH_MIN: '25'
        min_instances: 1
        max_instances: 3
  - scalability:
      type: tosca.policies.Scaling.MiCADO.VirtualMachine.CPU
      targets: [ worker-node ]
```

HUN-REN Cloud

```
tosca.policies.Scaling.MiCADO.Container.CPU.stressng:
  derived_from: tosca.policies.Scaling.MiCADO
  description: base MiCADO policy defining data sources, constants, queries,
  properties:
    alerts:
      type: list
      description: pre-define alerts for container CPU
      default:
      - alert: service_overloaded
        expr: 'avg(rate(container_cpu_usage_seconds_total{container_label_io_
        for: 30s
      - alert: service_underloaded
        expr: 'avg(rate(container_cpu_usage_seconds_total{container_label_io_
        for: 30s
      required: true
    scaling_rule:
      type: string
      description: pre-define scaling rule for container CPU
      default: |
        if len(m_nodes) == m_node_count:
          if service_overloaded and m_node_count > m_container_count:
            m_container_count+=1
          if service_underloaded:
            m_container_count-=1
        else:
          print('Transient phase, skipping update of containers...')
      required: true
```

Custom Python code

```
tosca.policies.Scaling.MiCADO.VirtualMachine.CPU.stressng:
  derived_from: tosca.policies.Scaling.MiCADO
  description: base MiCADO policy defining data sources, constants, queries, alert
  properties:
    alerts:
      type: list
      description: pre-define alerts for VM CPU
      default:
      - alert: node_overloaded
        expr: '(100-(avg(rate(node_cpu_seconds_total{node="{{ NODE_NAME }}", mode=
        for: 1m
      - alert: node_underloaded
        expr: '(100-(avg(rate(node_cpu_seconds_total{node="{{ NODE_NAME }}", mode=
        for: 1m
      required: true
    scaling_rule:
      type: string
      description: pre-define scaling rule for VM CPU
      default: |
        if len(m_nodes) <= m_node_count and m_time_since_node_count_changed > 60:
          if node_overloaded:
            m_node_count+=1
          if node_underloaded:
            m_node_count-=1
        else:
          print('Transient phase, skipping update of nodes...')
      required: true
```

Support for a large variety of clouds

Advanced security features

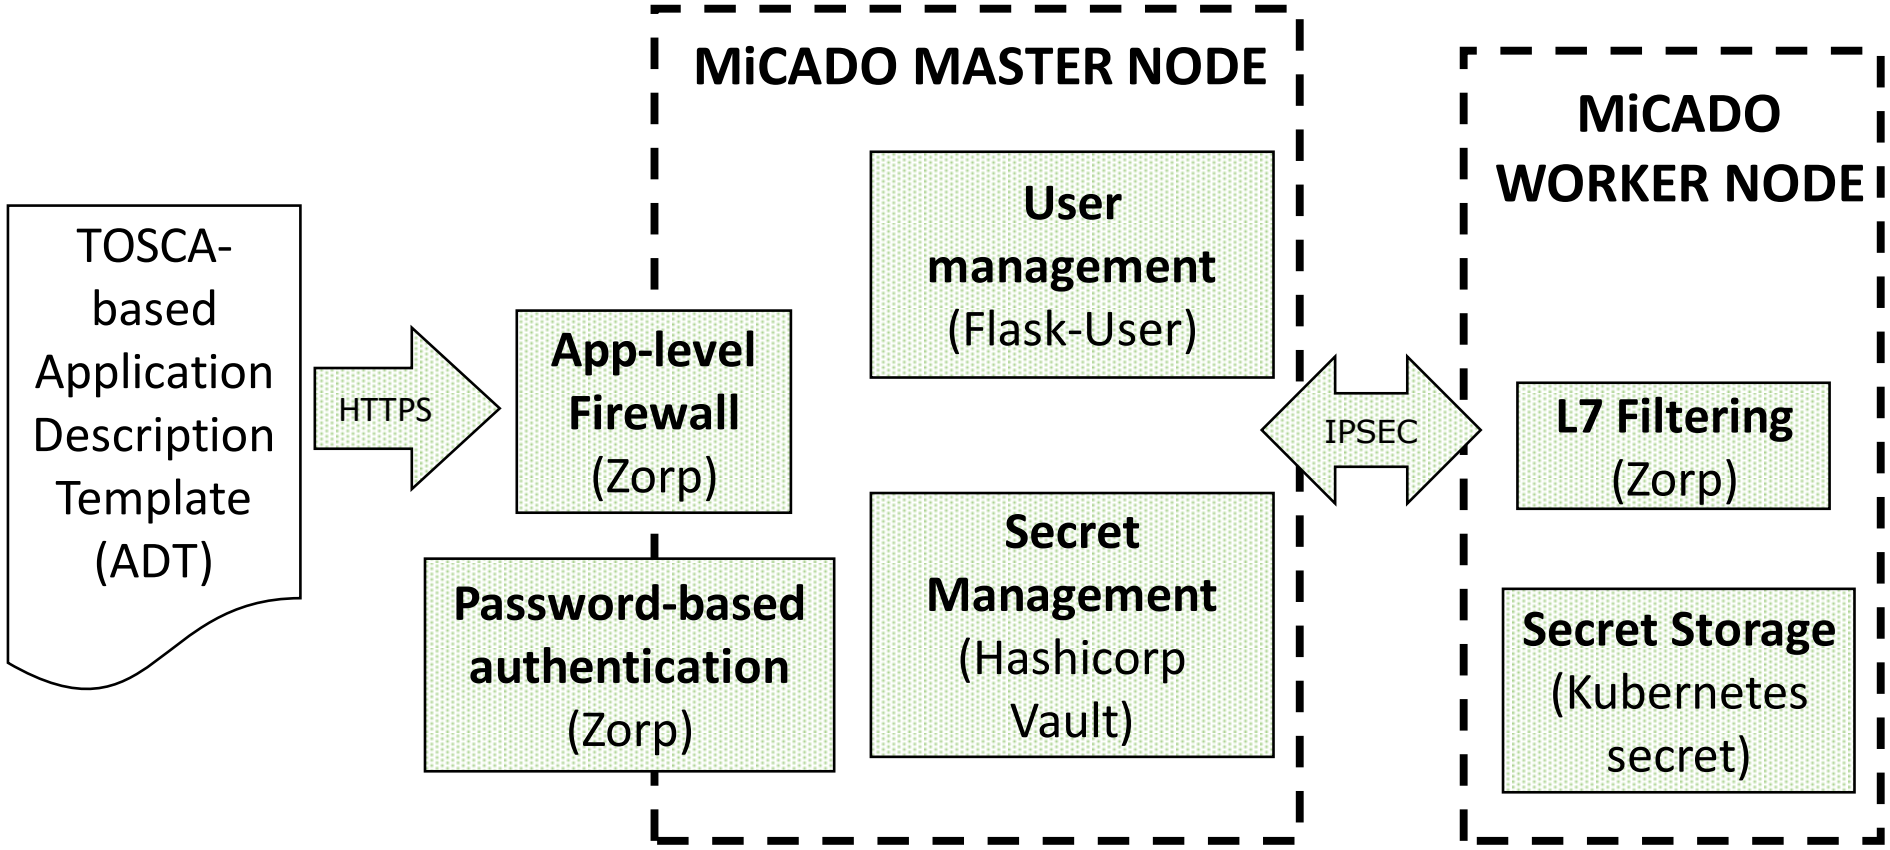

- implements **industry-standard** best practices
- provides **security functions lacking** in most cloud environments
- minimize the need of **user-supplied configuration**
- **pluggable** architecture
- validated by **penetration** testing

MiCADO and job execution

- Large number of jobs results in significant overall execution time
- Usually Restricted to complete all jobs by a deadline
  - Where to put the jobs?
  - How to distribute?
  - How to execute (in containers)
  - How to liaise with deadline?

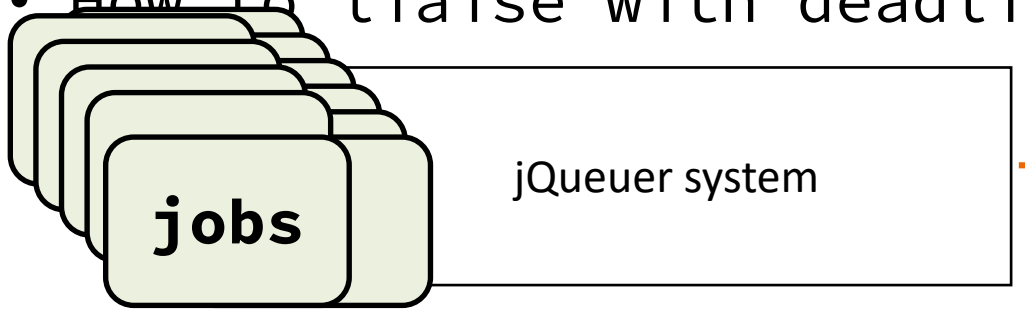

ADT: infrastructure and scaling rules

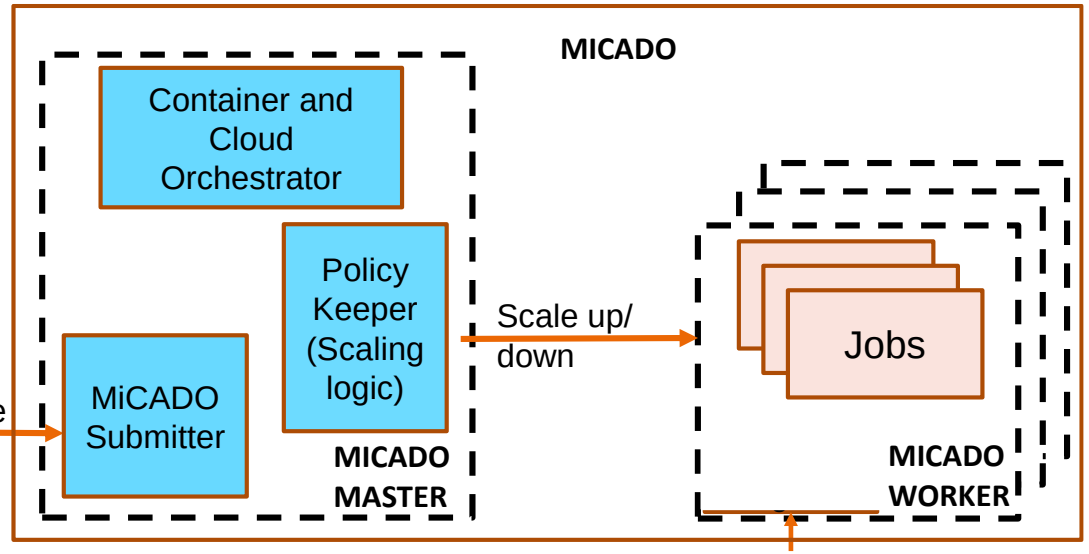

MiCADO and job execution

```
{
  "container_name": "worker",
  "single_task_duration": 60,
  "experiment_deadline": 1800,
  "jobs":
  [
    Job 1:
    {
      Task 1:
      ...
      Task N:
    }
    ...
    Job N:
  ]
}
```

Experiment wide global parameters

List of tasks

List of jobs

Experiment definition

Algorithm 1: Scaling logic

Result: Calculate the number of required Worker nodes at any point in time

```
items = TOTALJOBS - (COMPLETEDJOBS + FAILEDJOBS);
if COMPLETEDJOBS > 0 then
    aet = average execution time of all completed jobs;
else
    aet = user specified estimated single job execution time;
end
rt = remaining time to deadline;
N = ceil(aet/((rt - aet * 0.20)/items));
if (N ≤ 0) or (N ≥ items) then
    N = items
end
if N == current number of nodes then
    Do nothing
else
    Scale up/down accordingly
end
```

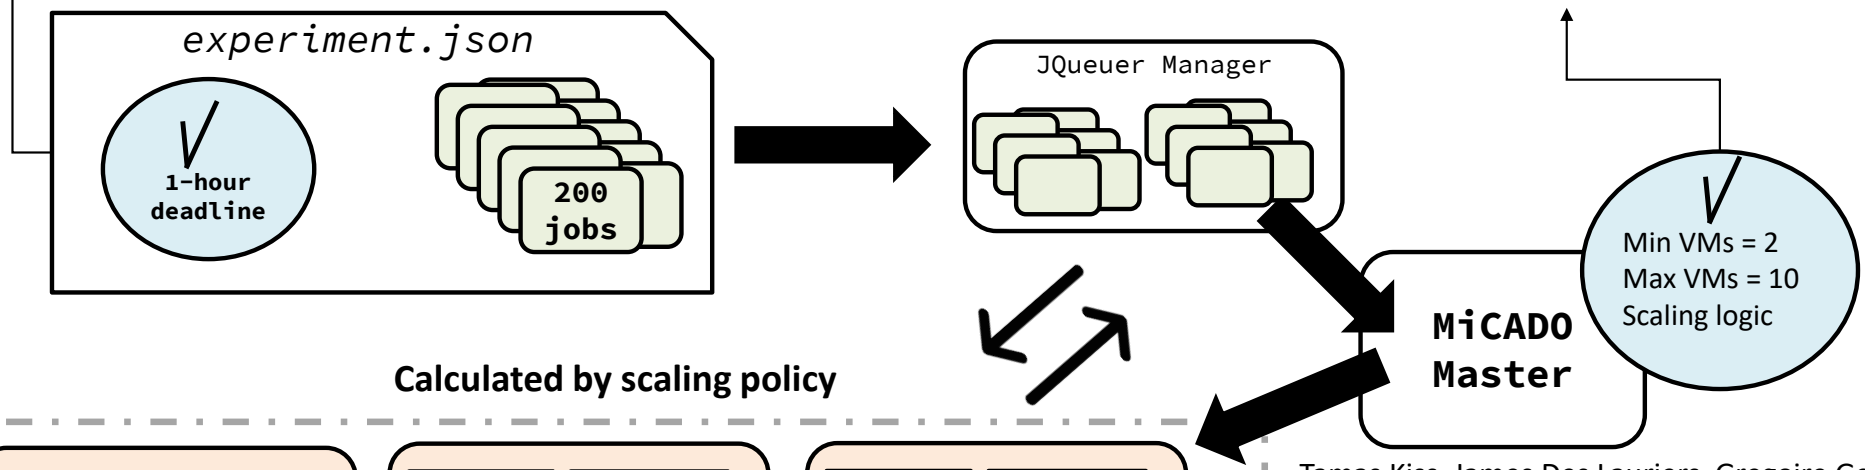

Tamas Kiss, James Des Lauriers, Gregoire Gesmier, Gabor Terstyanszky, Gabriele Pierantoni, Osama Abu Oun, Simon JE Taylor, Anastasia Anagnostou, Jozsef Kovacs, **Cloud-agnostic Queuing System to Support the Implementation of Deadline-based Application Execution Policies**, Future Generation Computer Systems, Elsevier, Vol 101, December, 2019, pp 99-111

MiCADO EDGE/FOG Extension

- Solution using KubeEdge
- Automated deployment of microservices extended to edge nodes
- Monitoring information collected from edge workers
- Scaling/reconfiguration policies extended towards edge

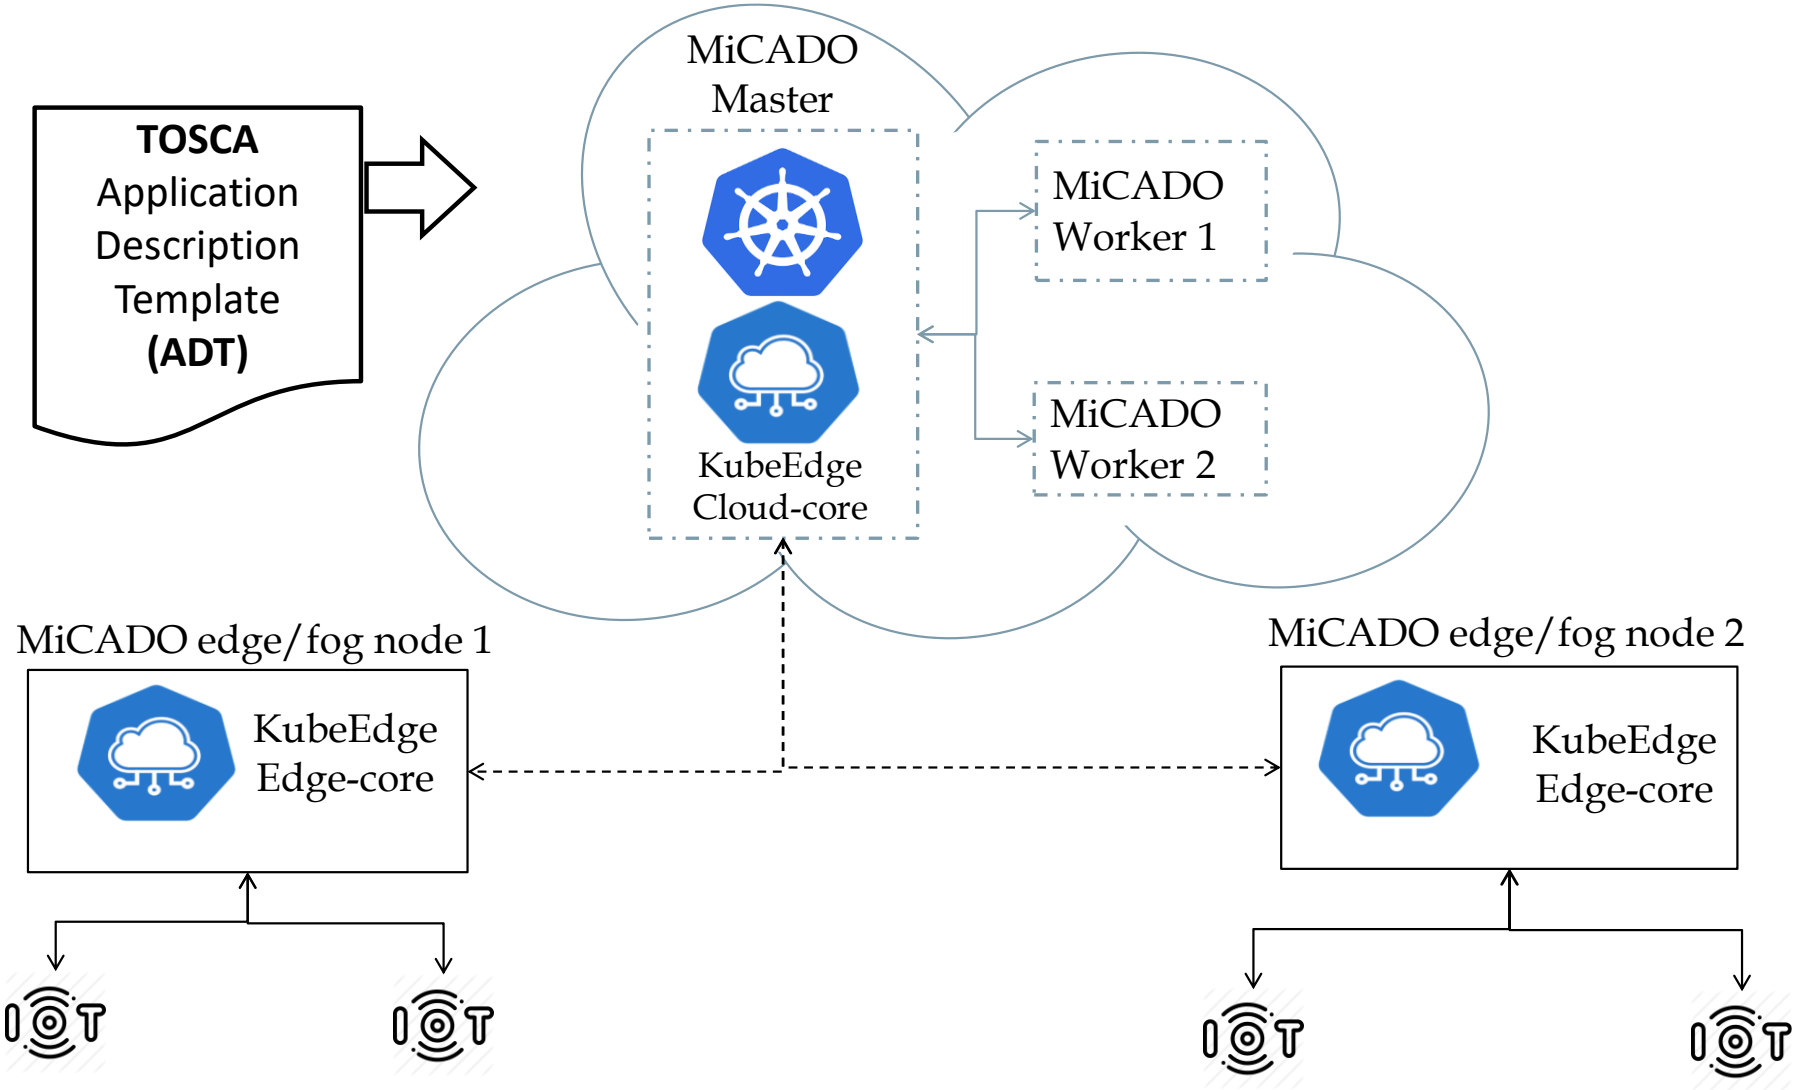

Architecture with Edge extension

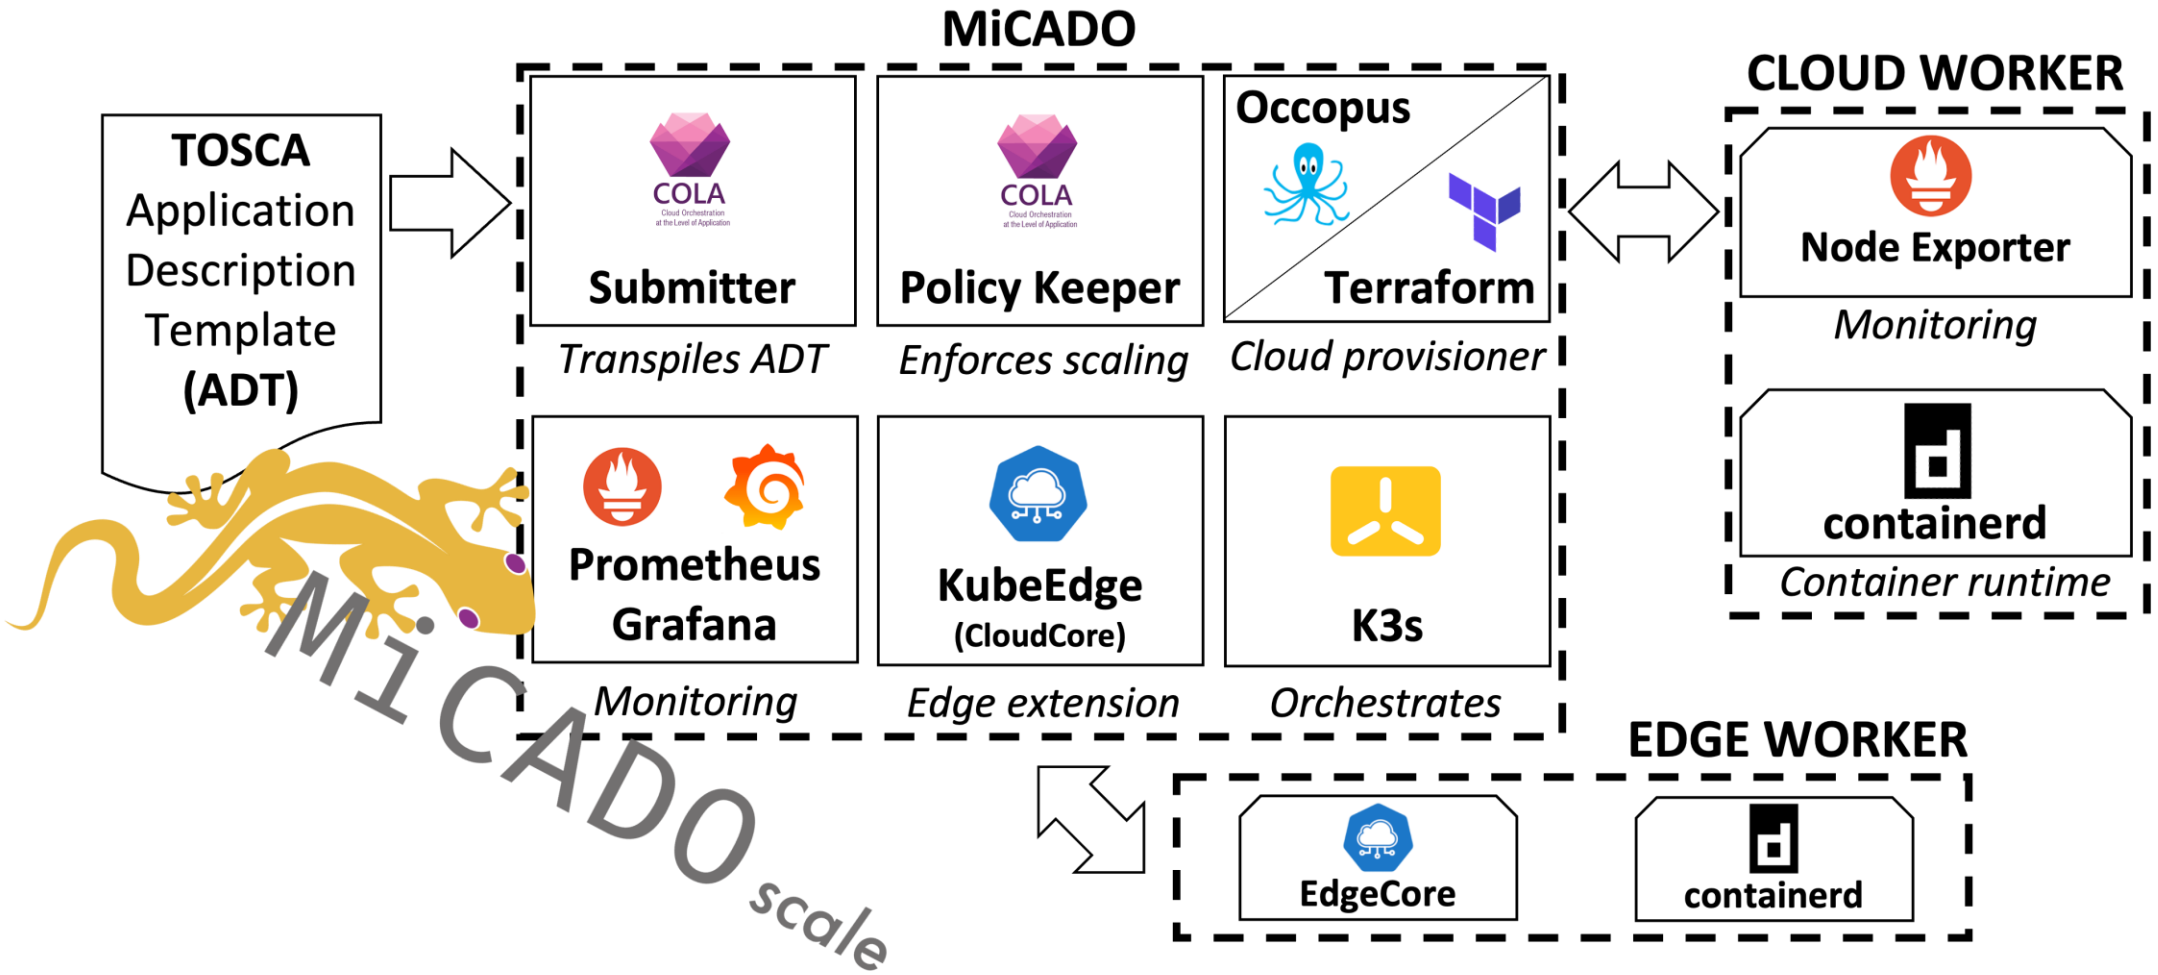

Now with a single uniform descriptor (ADT), the entire Cloud-to-Edge application can be described, E.g.

```
fogedge:
  type: tosca.nodes.MiCADO.Edge
  properties:
    public_ip: { get_input: ip_fog_node }
```

Edge Node ↑

auto-deployment of edge through ADT

```
fd-processor:
  type: tosca.nodes.MiCADO.Container.Application.Docker.Deployment
  properties:
    image: uowcpc/fd-edge-processor
    env:
    - name: SLEEP_PERIOD
      value: "2.0"
  requirements:
  - host: fogedge
  - volume: docker-edge-grey-images-host-vol
  - volume: docker-edge-images-host-vol
```

Service Container ↑

MiCADO edge/fog extension – face detection application

- Cloud server stores images with faces
- Fog node recognises faces in images
- Edge device captures video stream

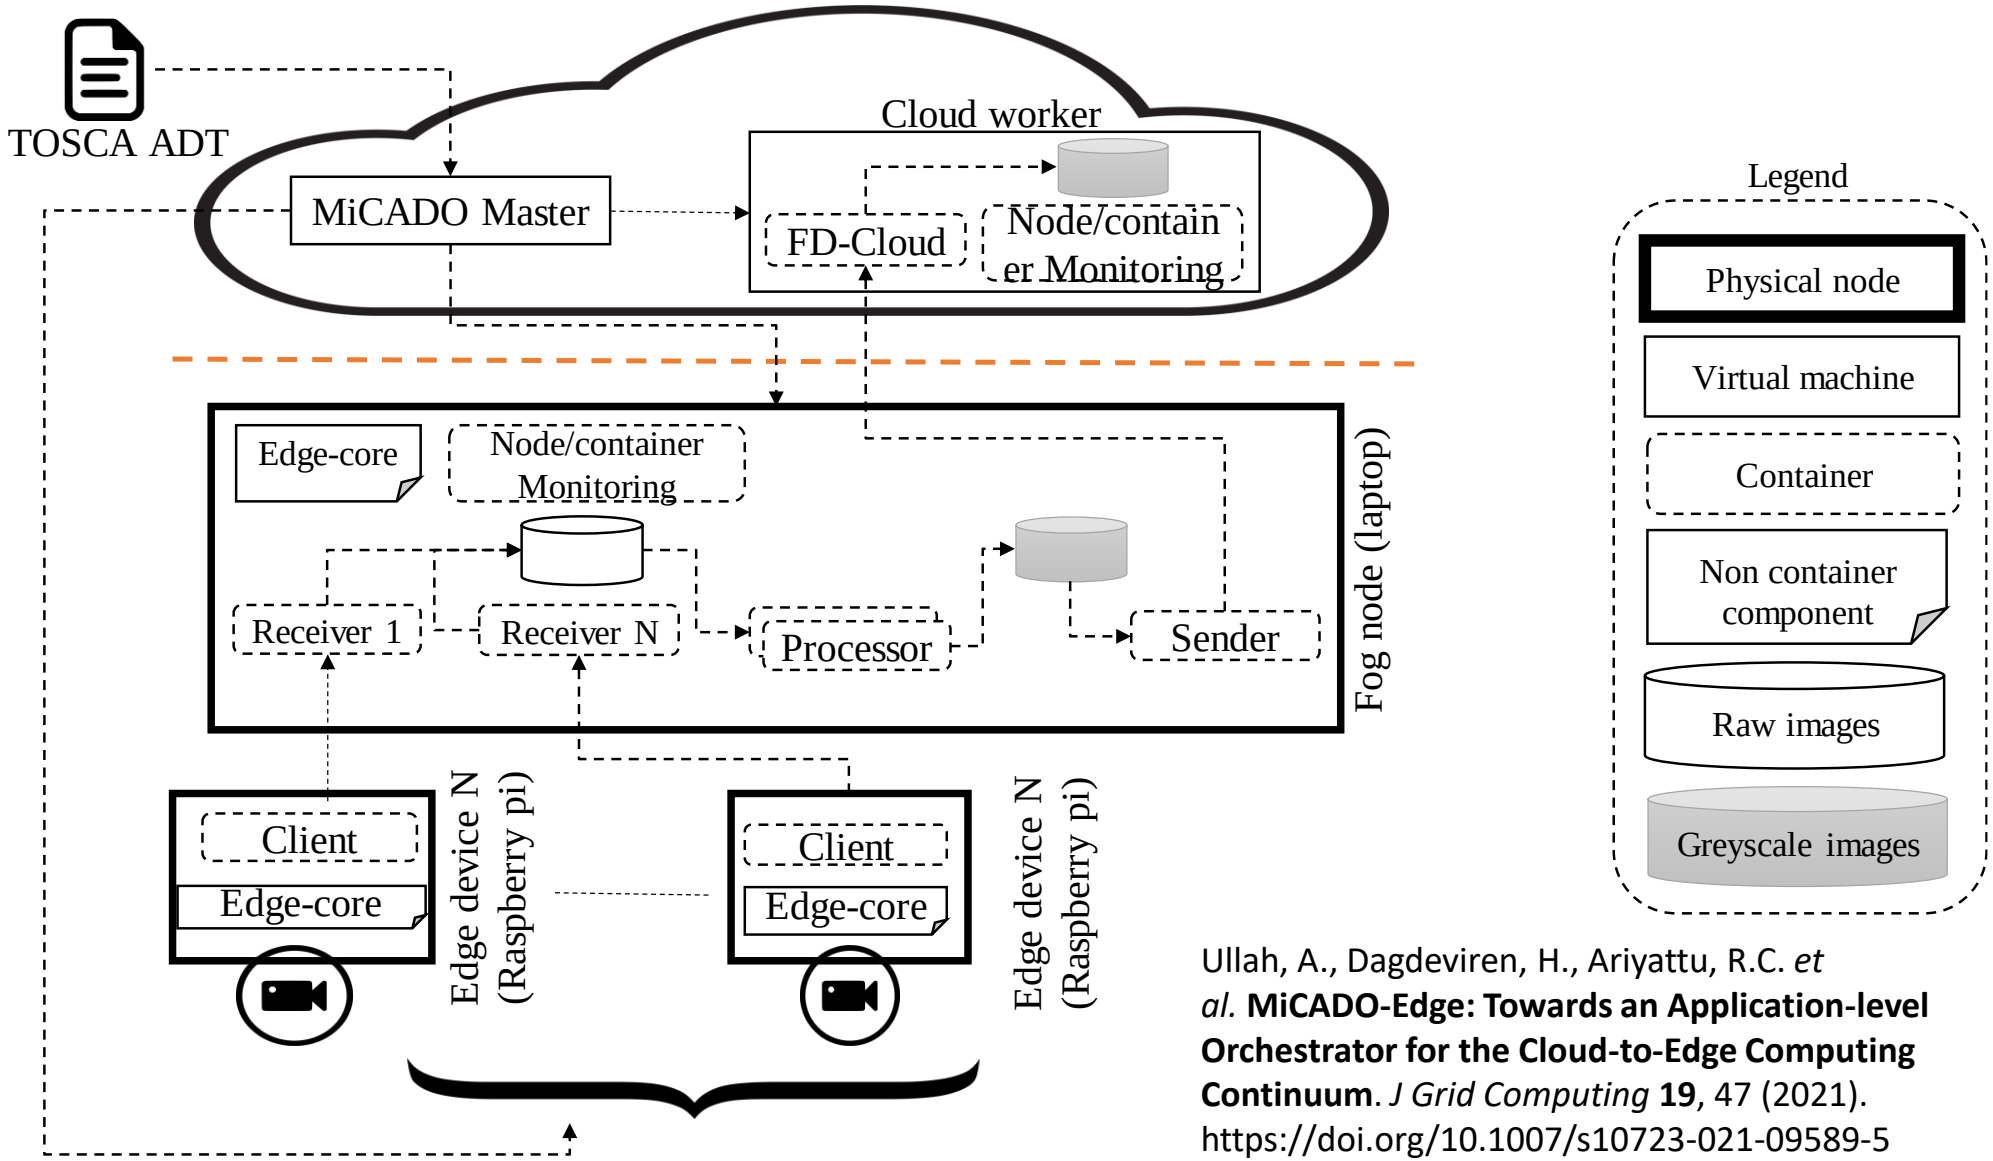

Ullah, A., Dagdeviren, H., Ariyattu, R.C. *et al.* **MiCADO-Edge: Towards an Application-level Orchestrator for the Cloud-to-Edge Computing Continuum**. *J Grid Computing* **19**, 47 (2021). https://doi.org/10.1007/s10723-021-09589-5

Deployment: Step 1: install micado-client

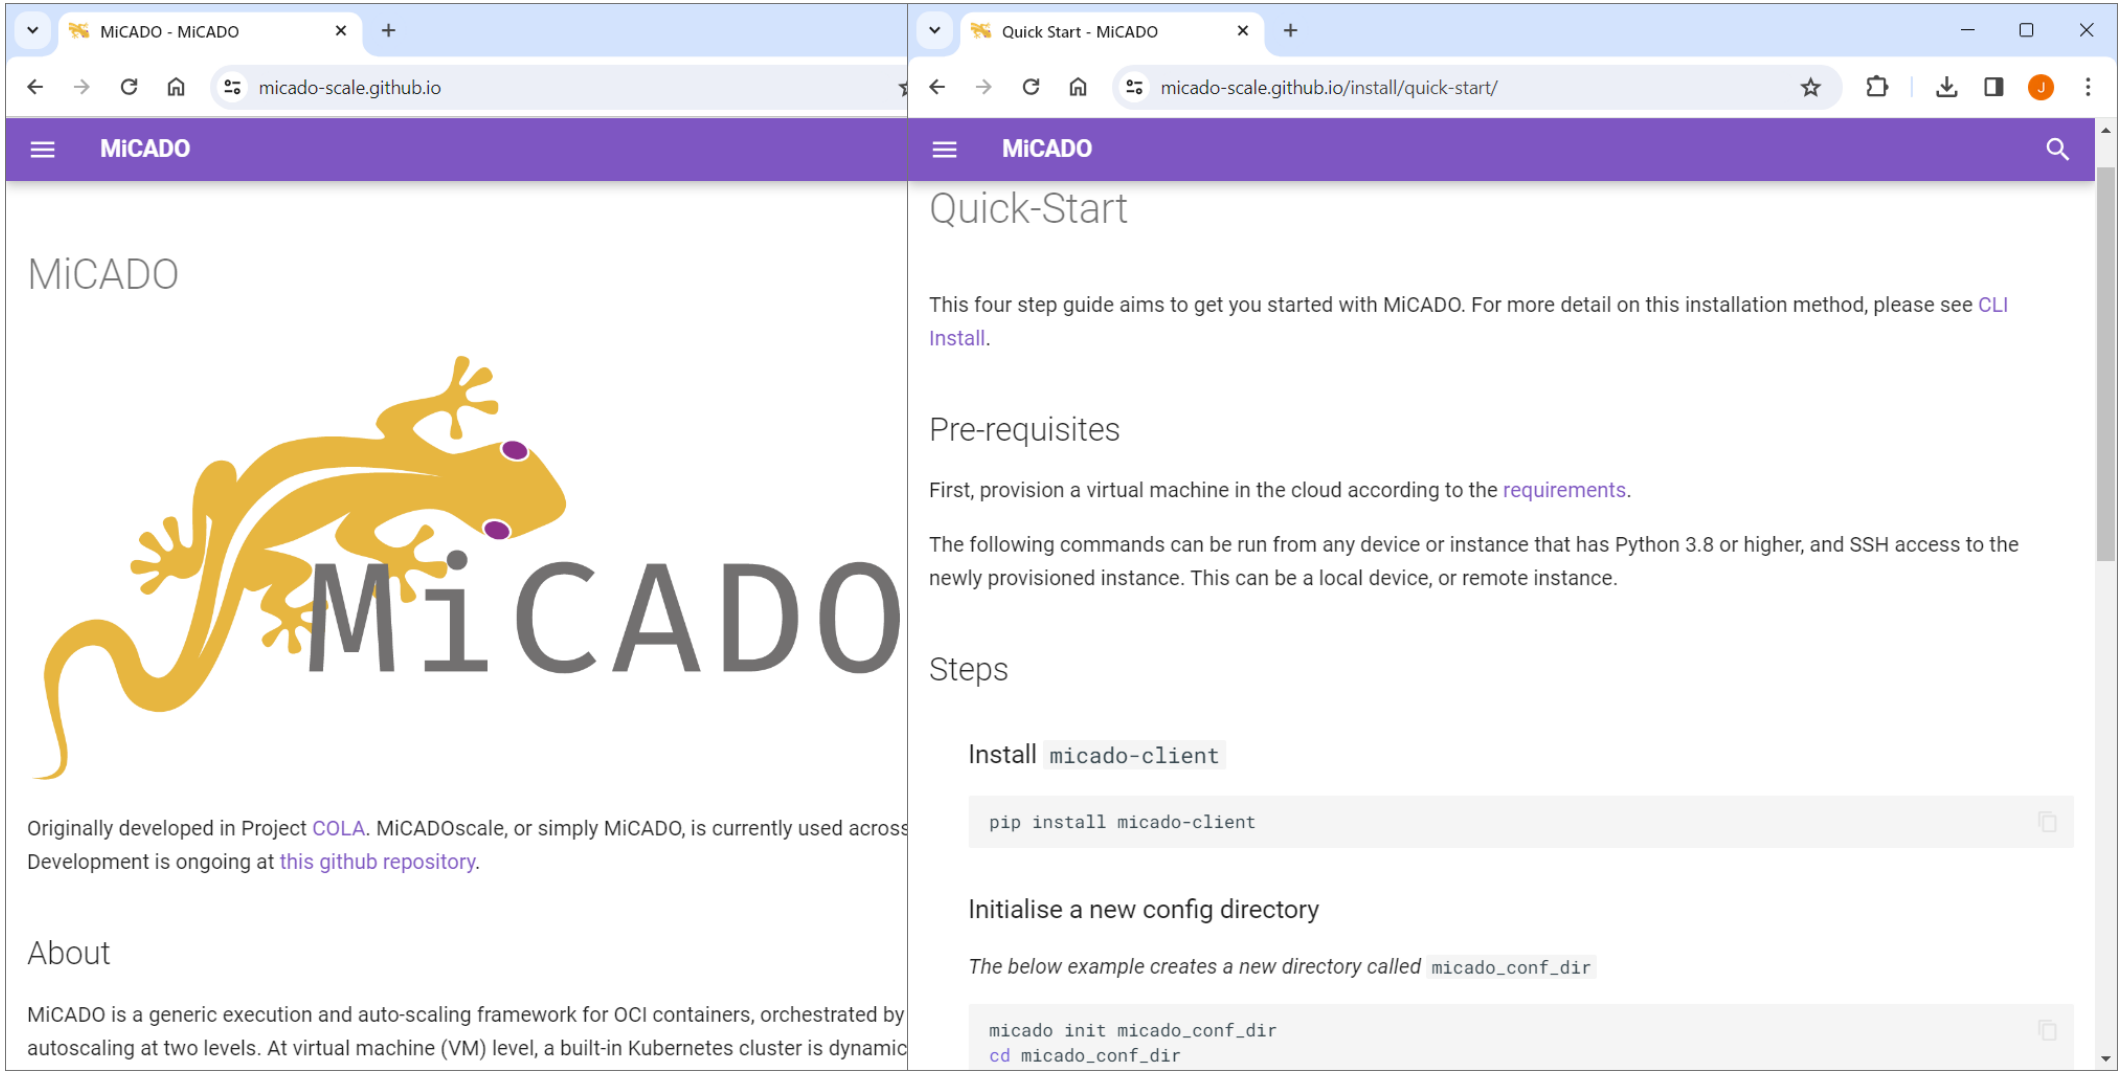

Step 2: launch a new VM for MiCADO

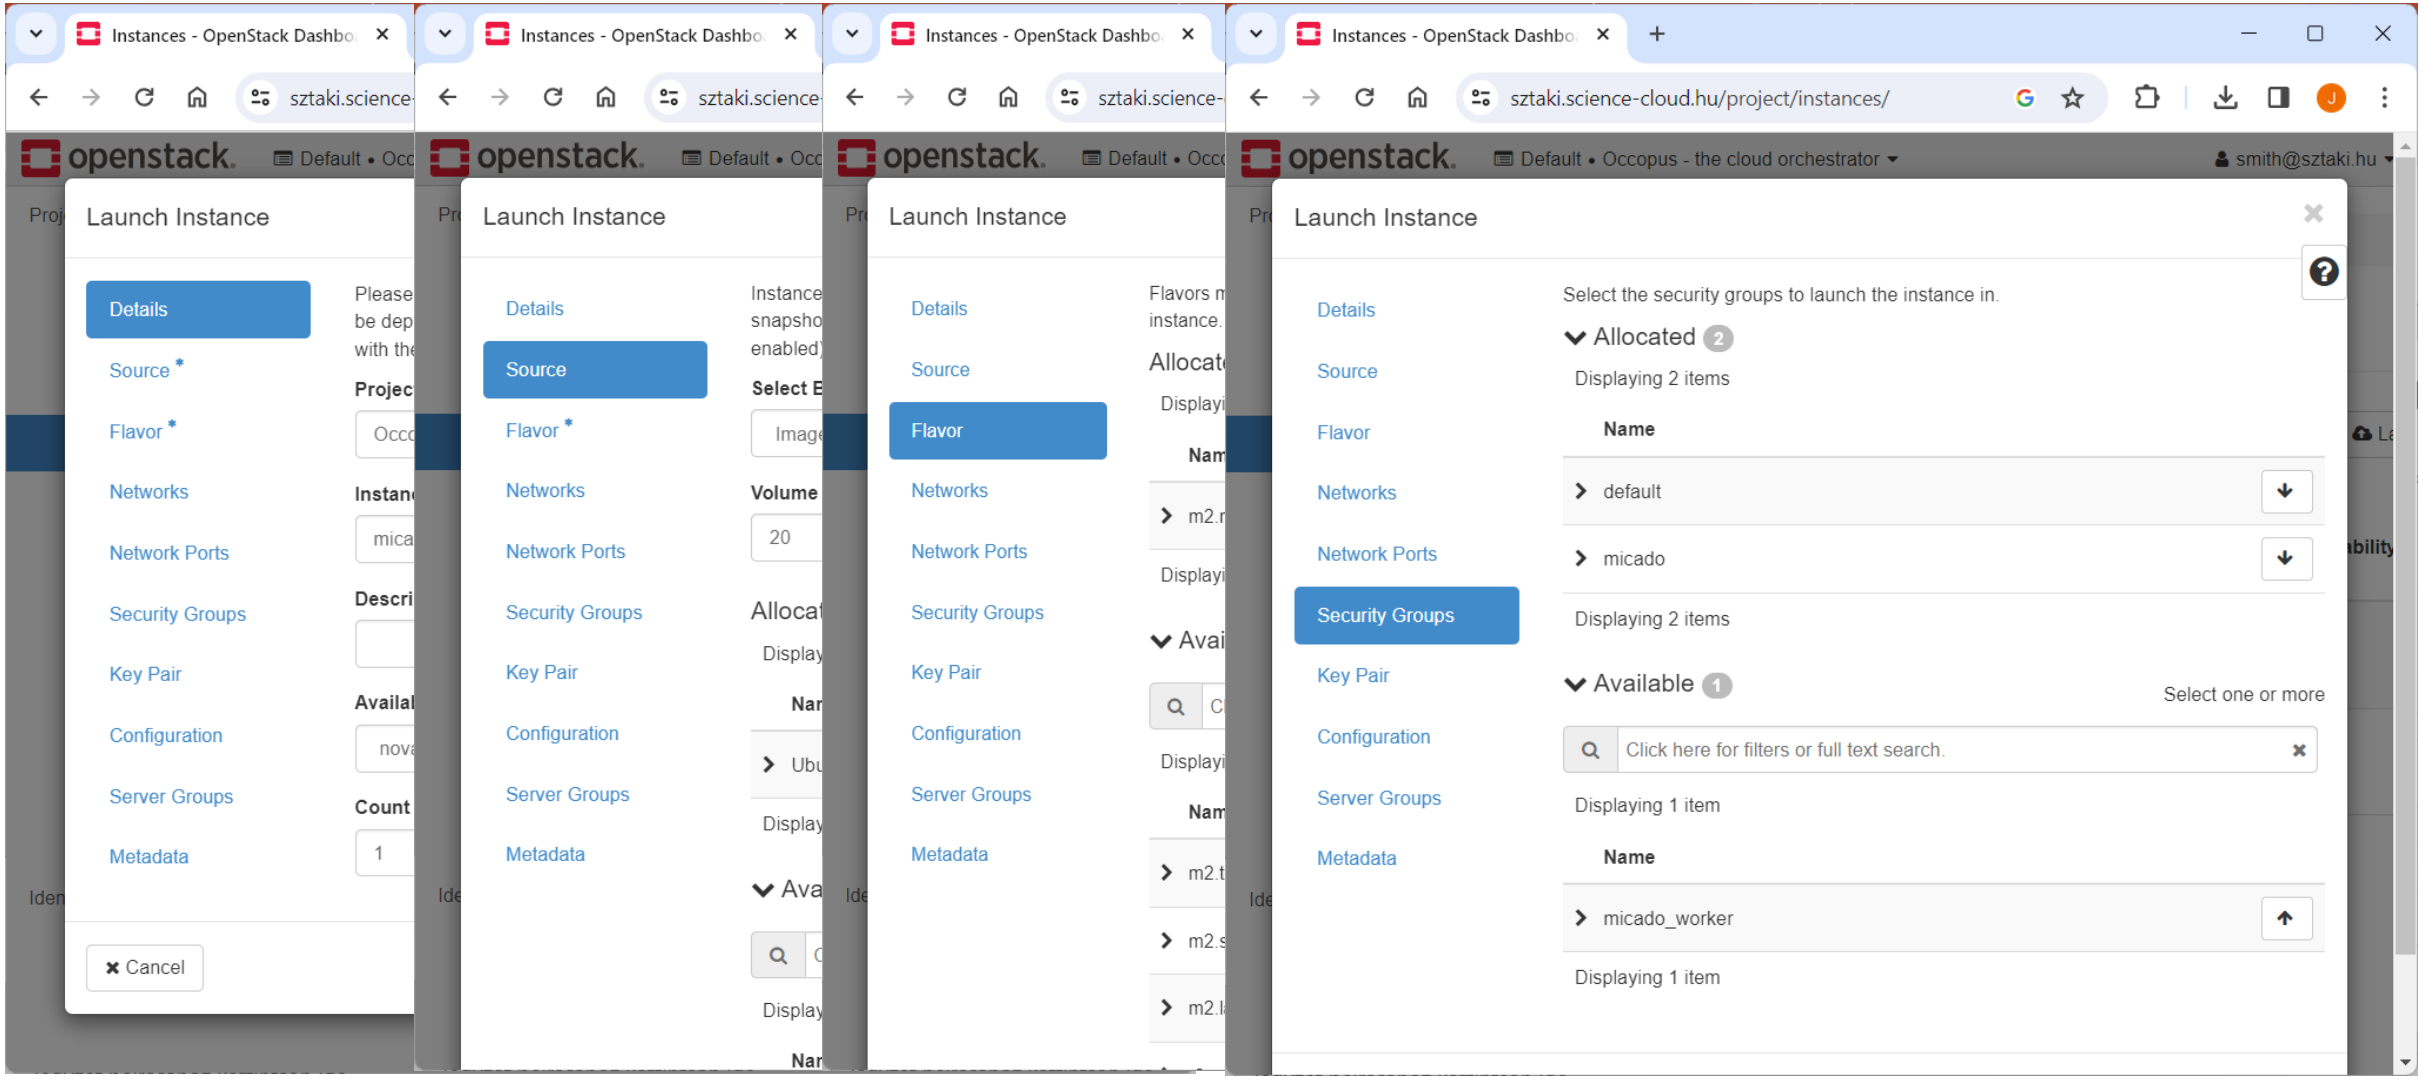

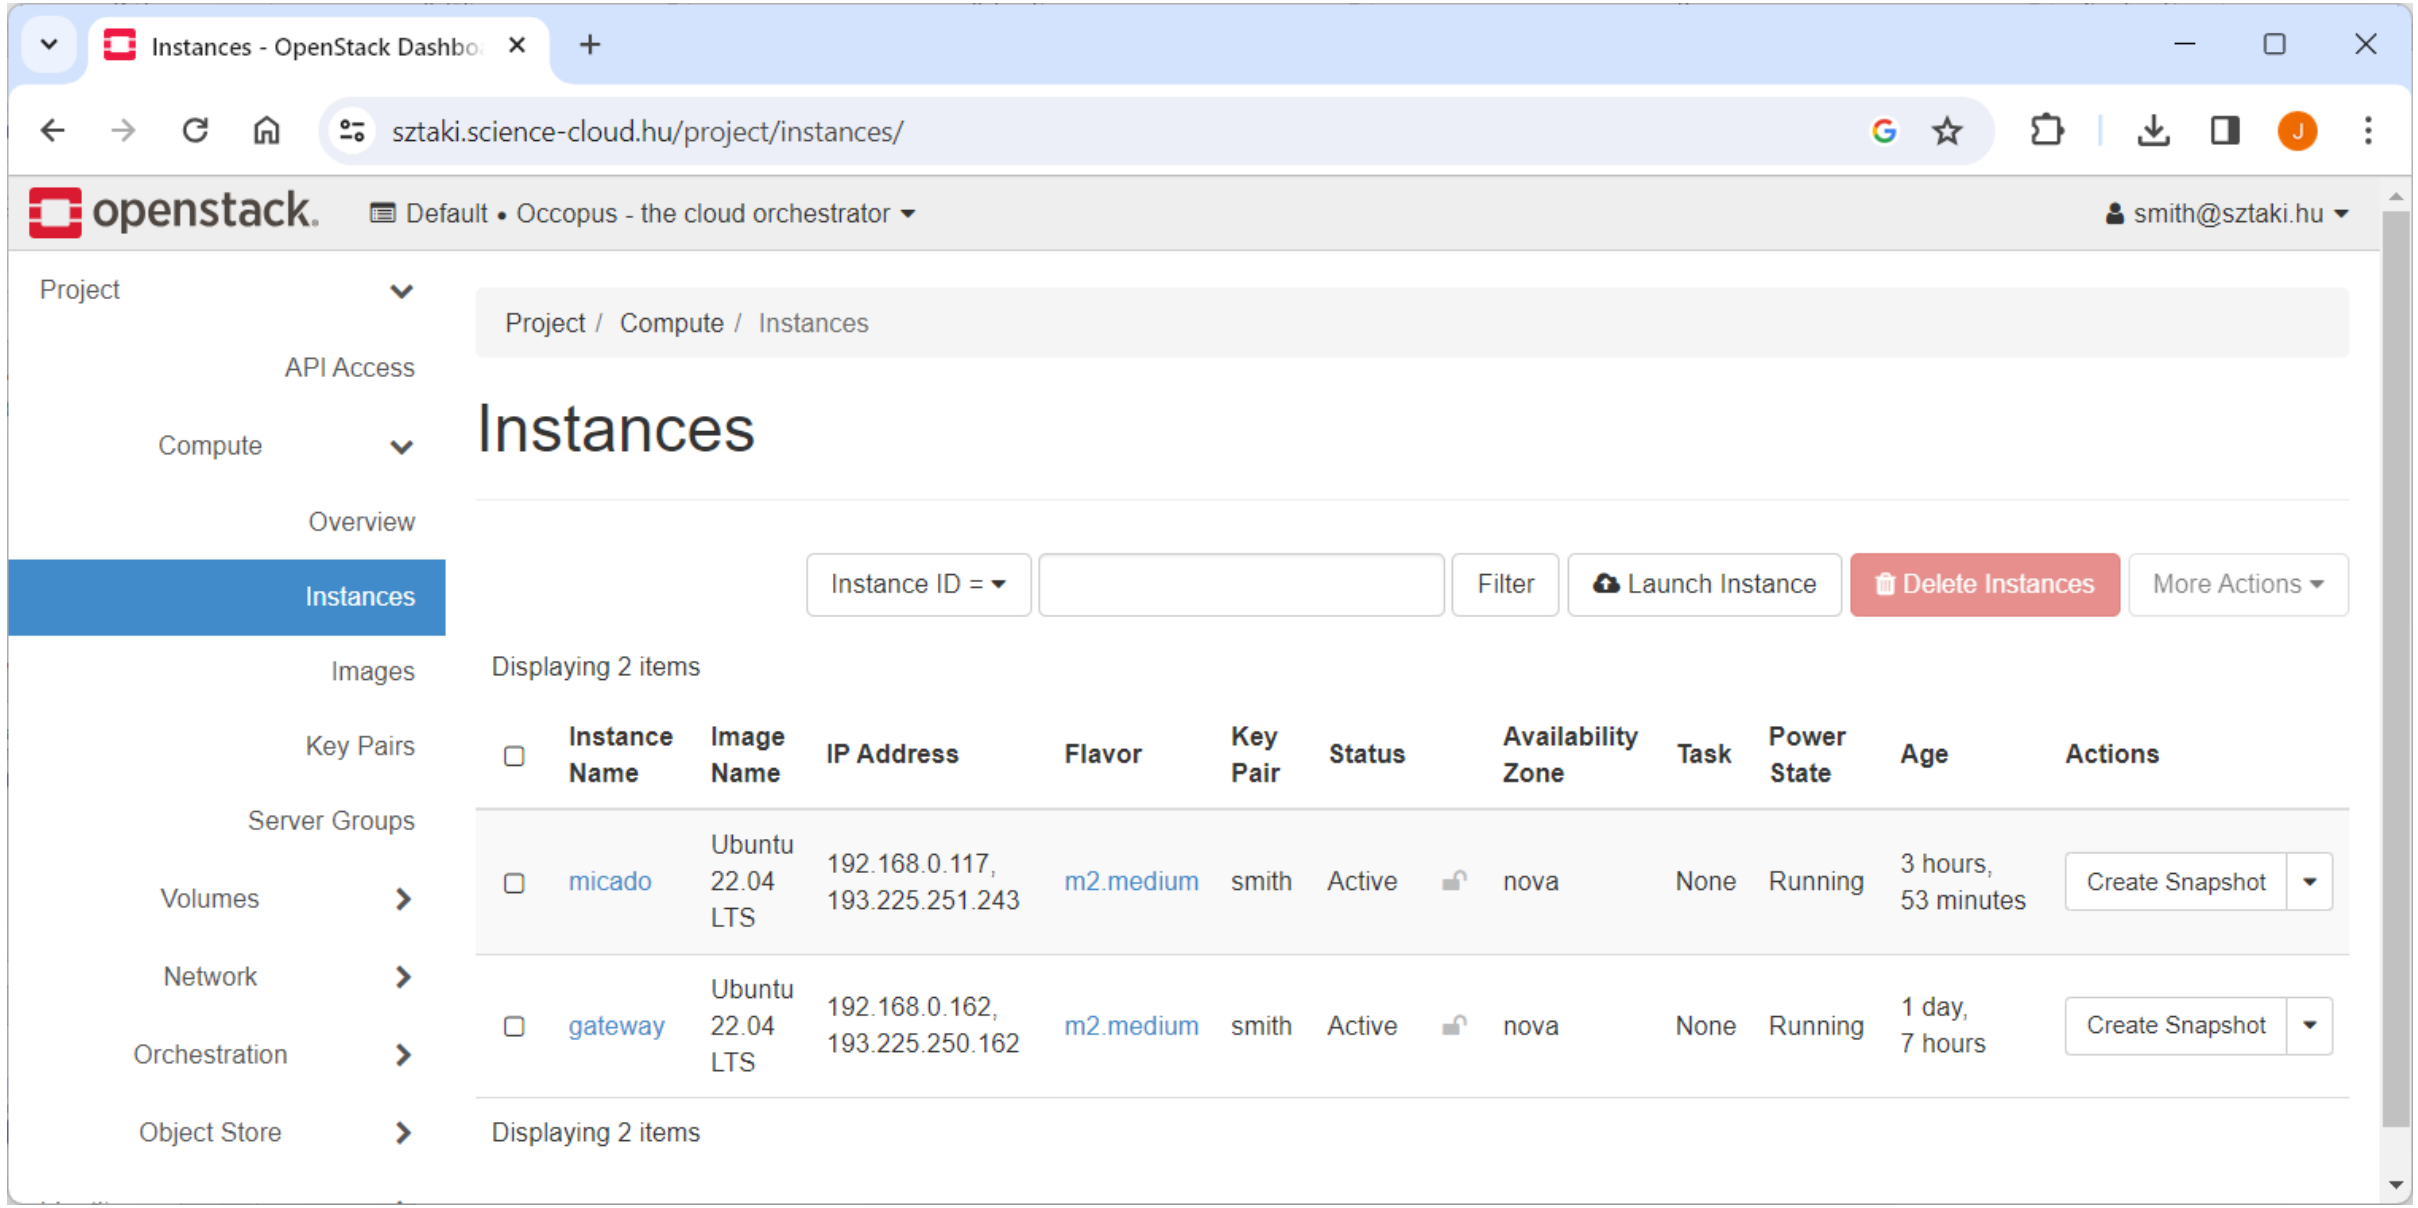
Instances - OpenStack Dashbo
sztaki.science-cloud.hu/project/instances/
openstack.
Default • Occopus - the cloud orchestrator
smith@sztaki.hu
Project
API Access
Compute
Overview
Instances
Images
Key Pairs
Server Groups
Volumes
Network
Orchestration
Object Store
Project / Compute / Instances
Instances
Instance ID =
Filter
Launch Instance
Delete Instances
More Actions
Displaying 2 items
Instance Name
Image Name
IP Address
Flavor
Key Pair
Status
Availability Zone
Task
Power State
Age
Actions
micado
Ubuntu 22.04 LTS
192.168.0.117, 193.225.251.243
m2.medium
smith
Active
nova
None
Running
3 hours, 53 minutes
Create Snapshot
gateway
Ubuntu 22.04 LTS
192.168.0.162, 193.225.250.162
m2.medium
smith
Active
nova
None
Running
1 day, 7 hours
Create Snapshot
Displaying 2 items

Step 3: configure 5 groups of details

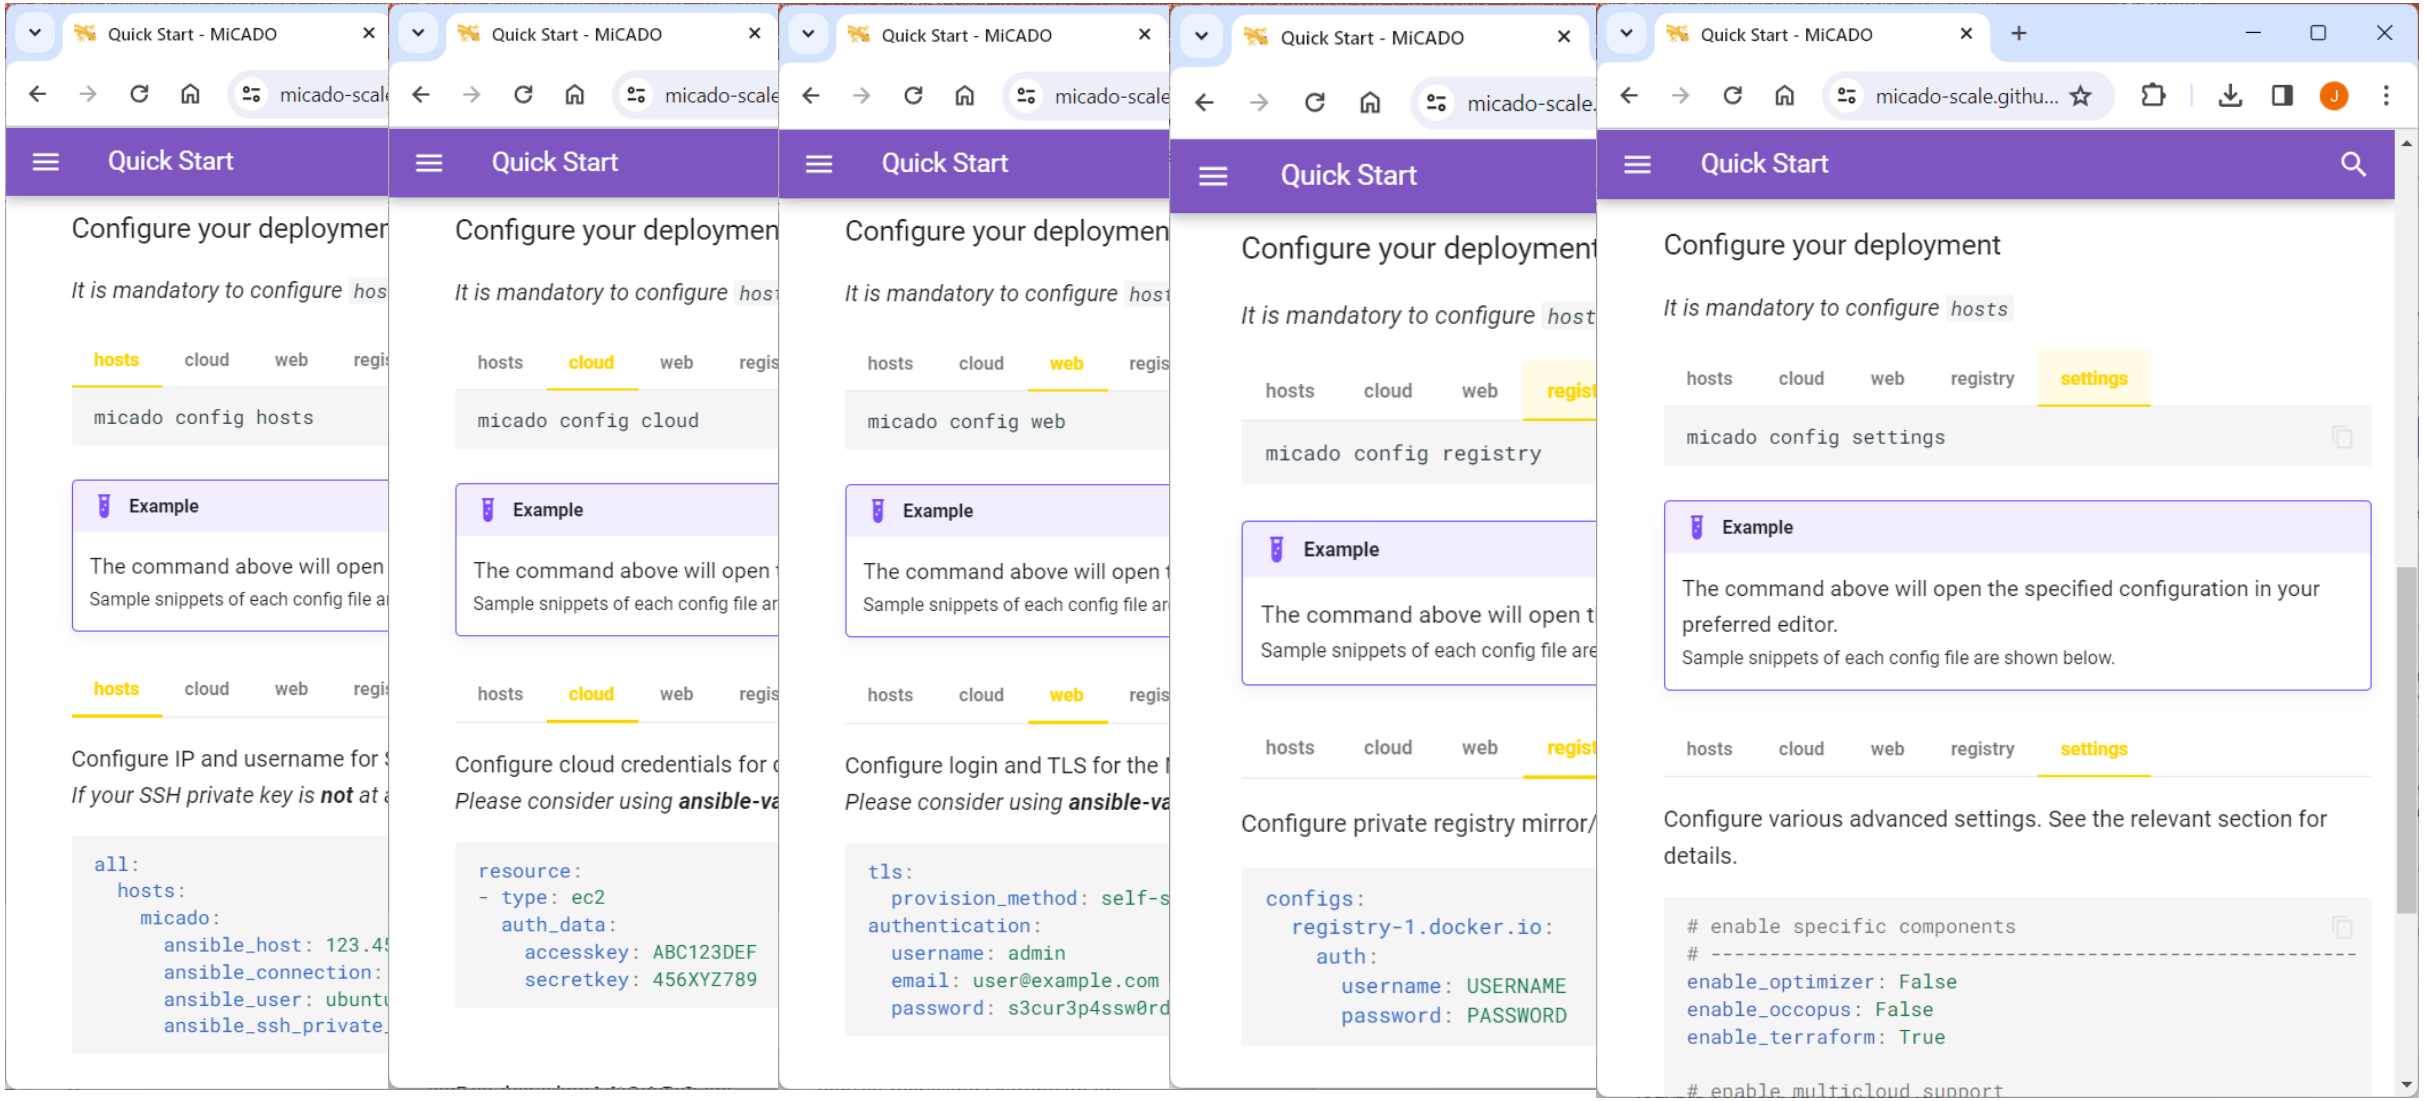

Step 4: deploy micado

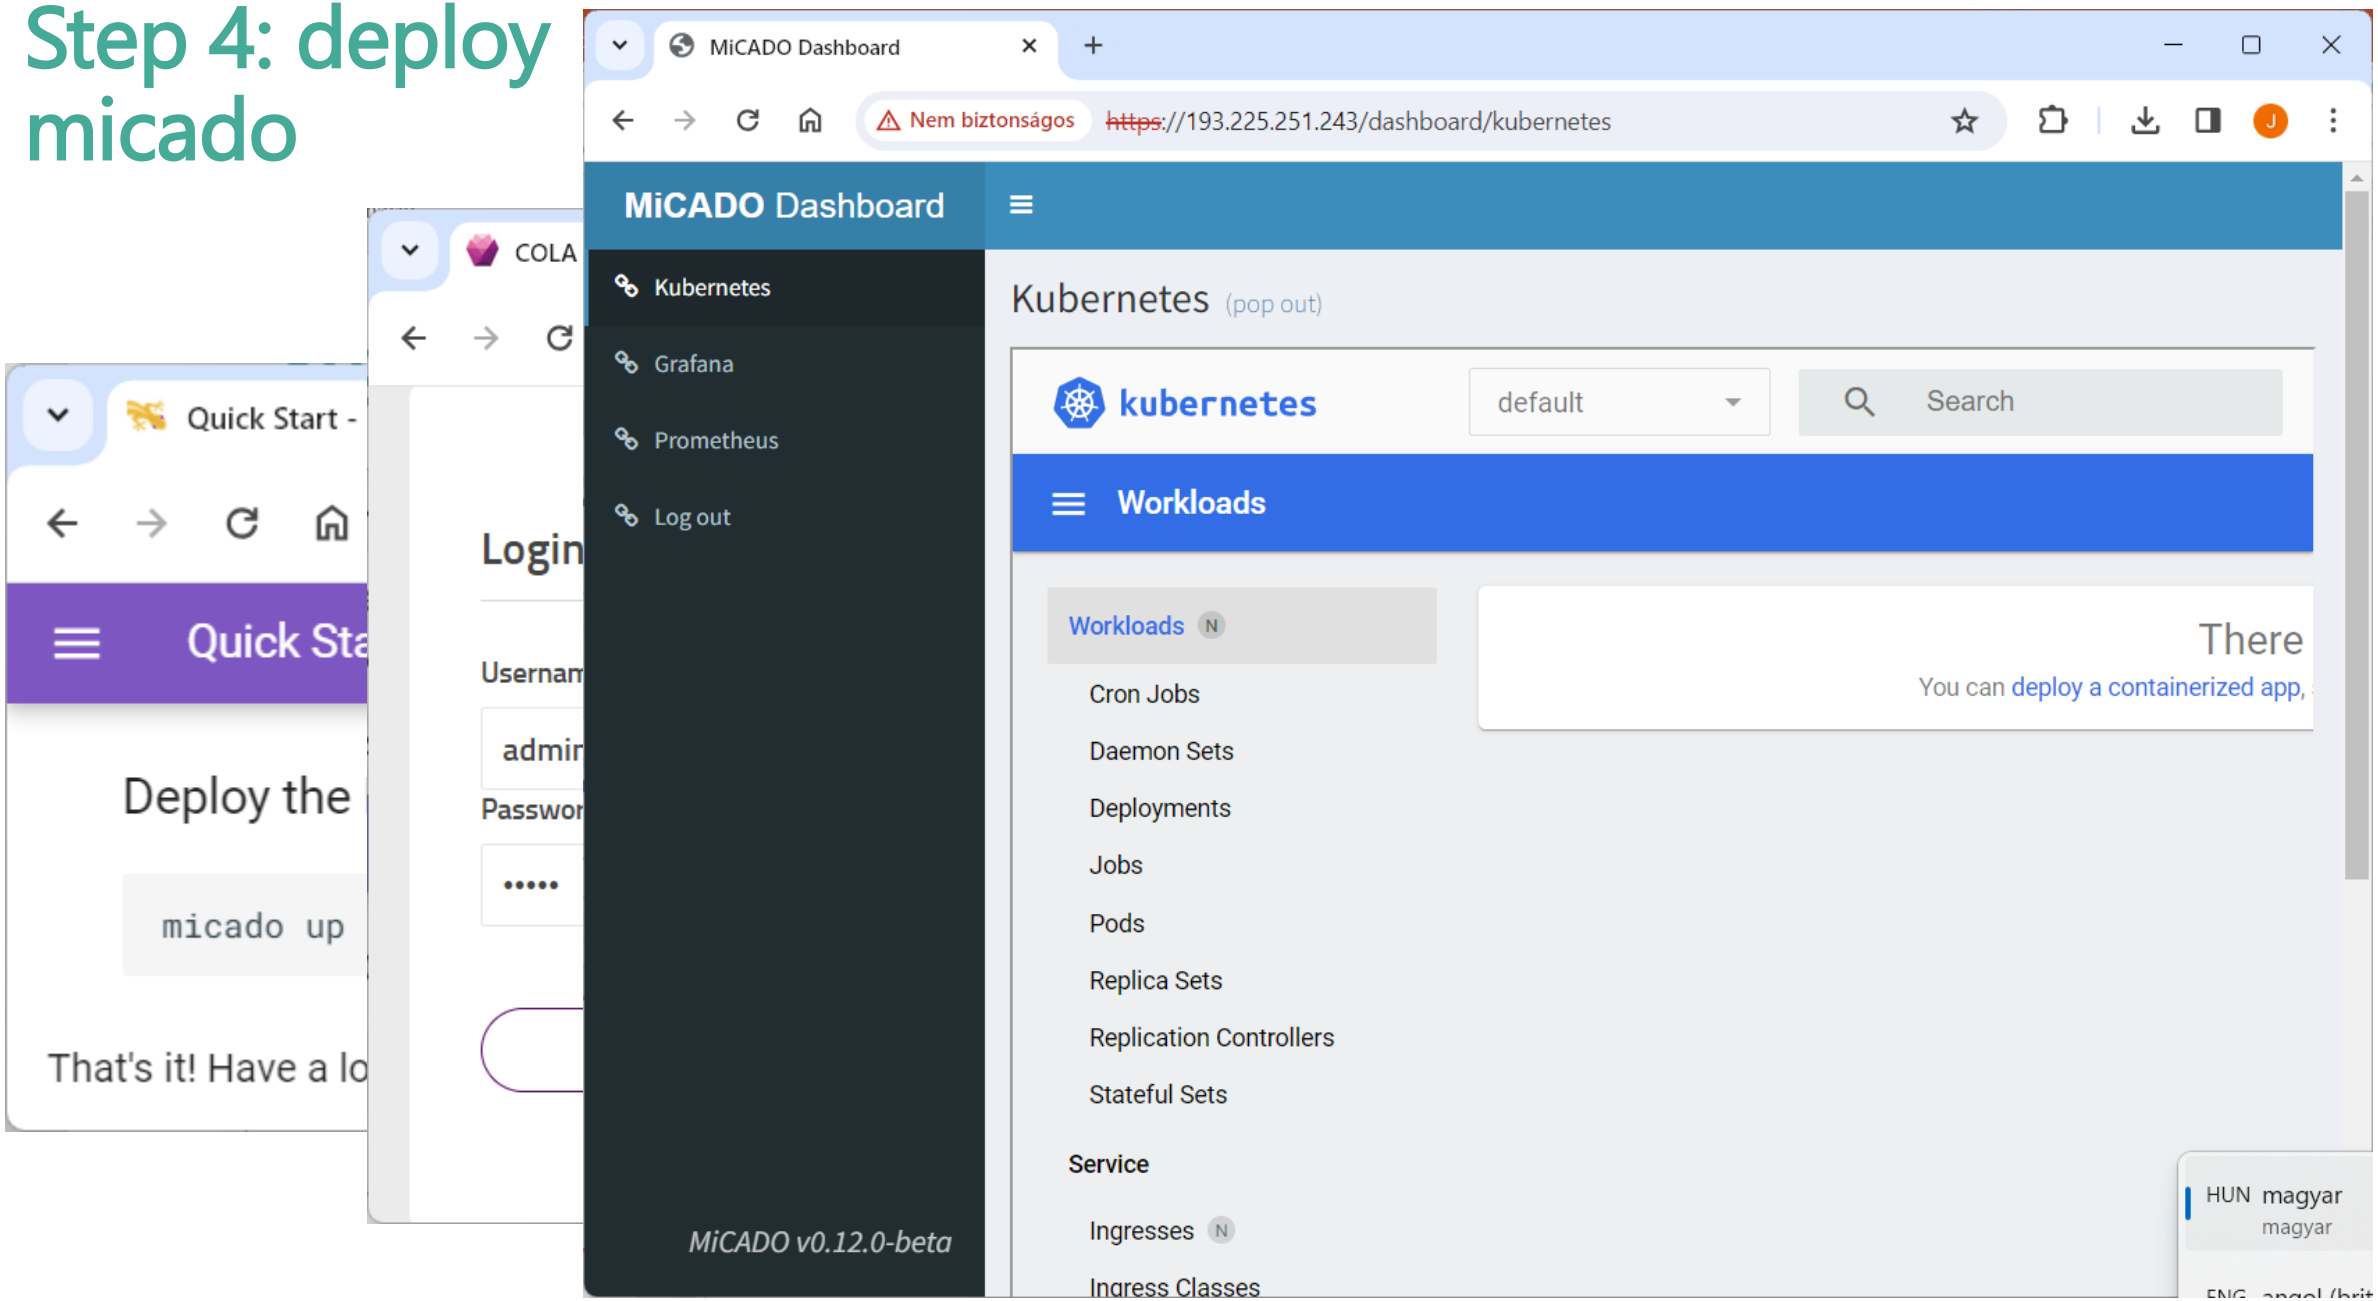

Step 5: try demos

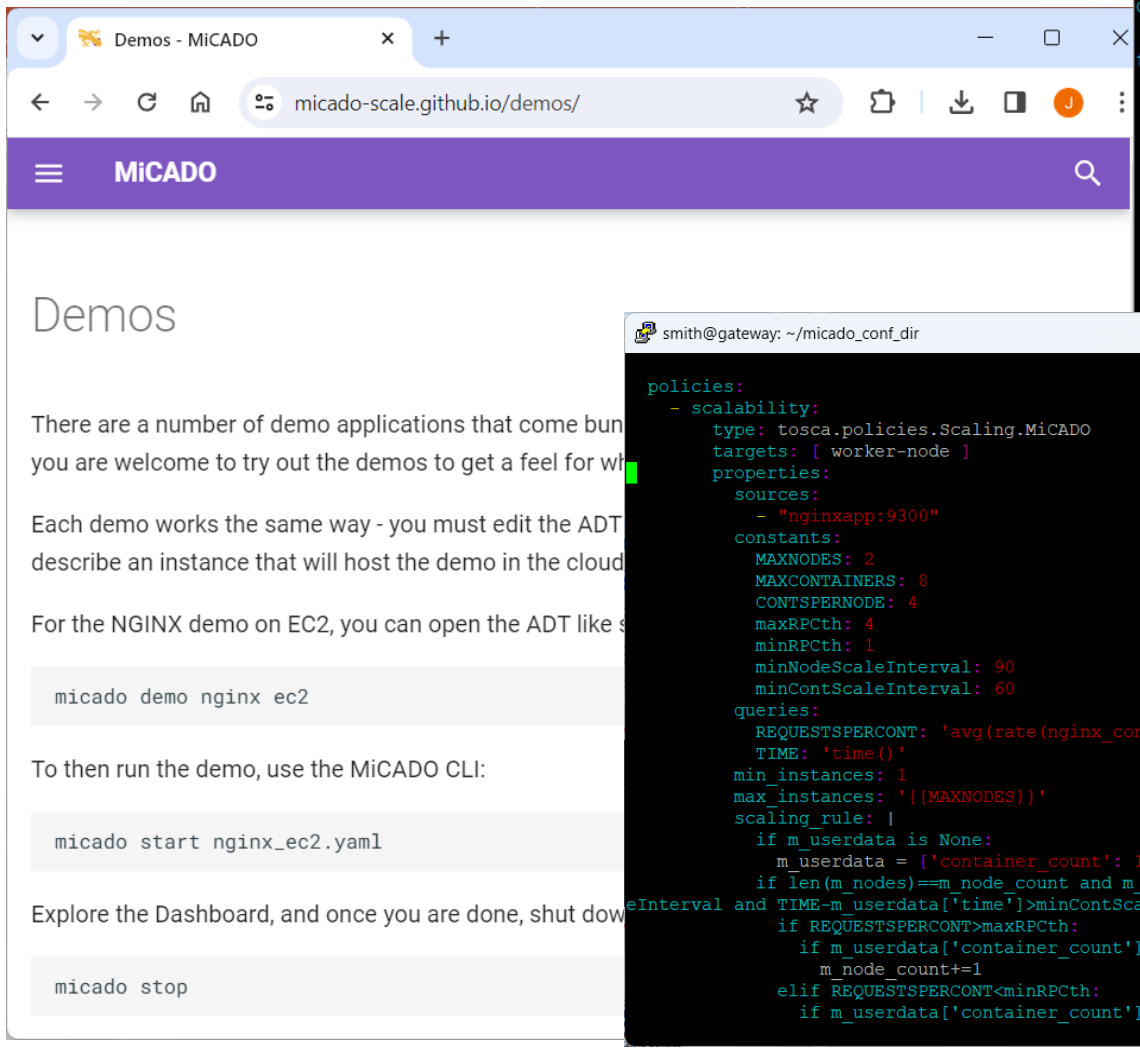

Step 6: build your own application

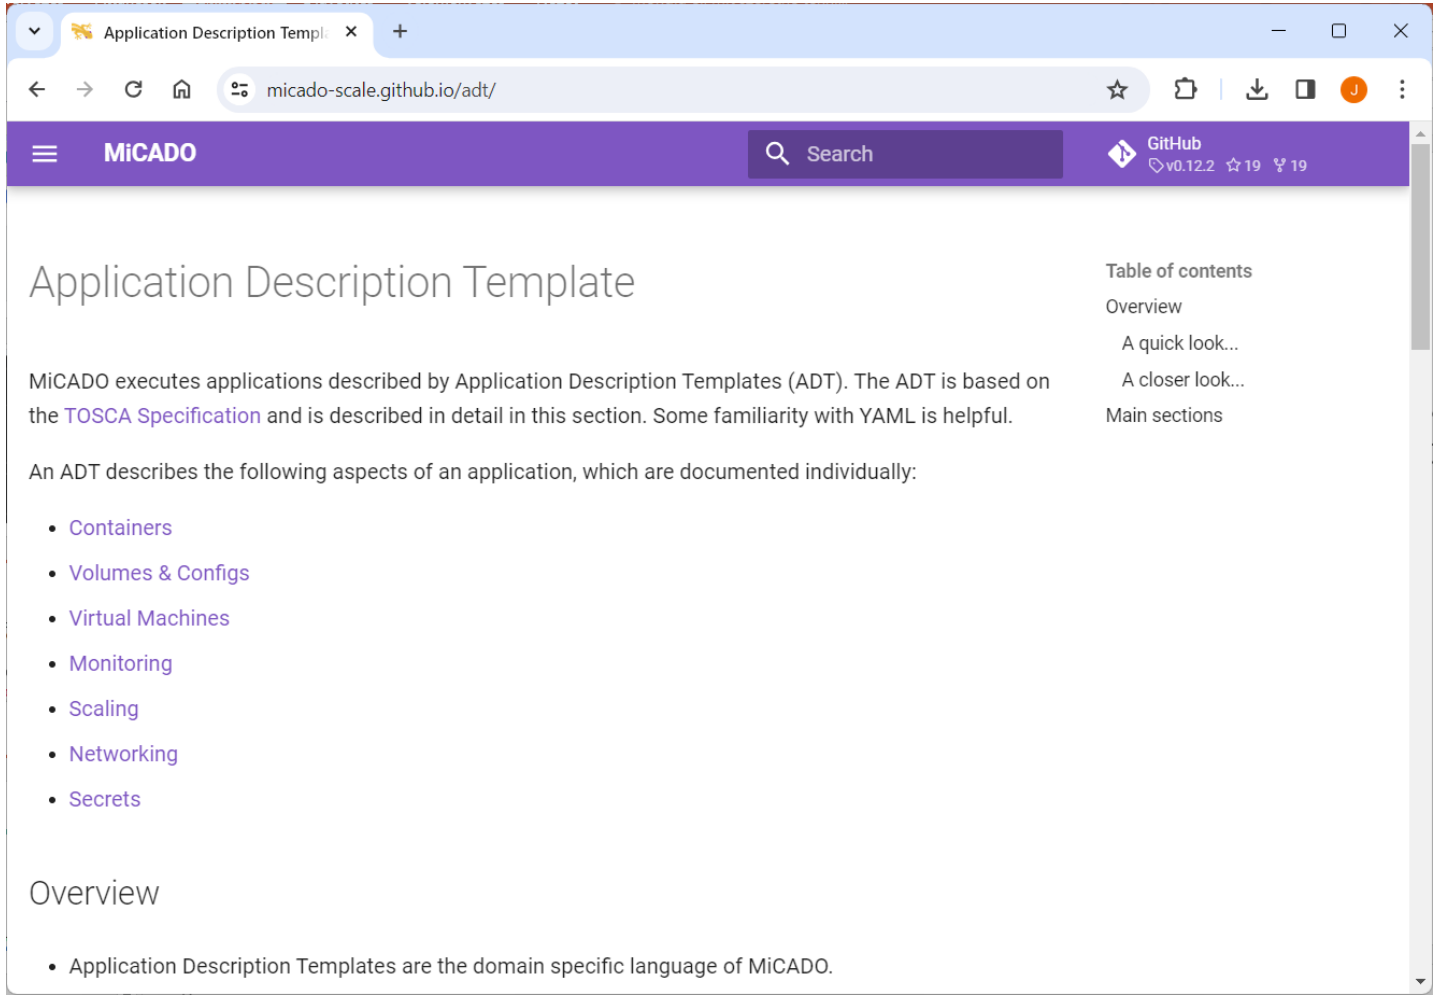

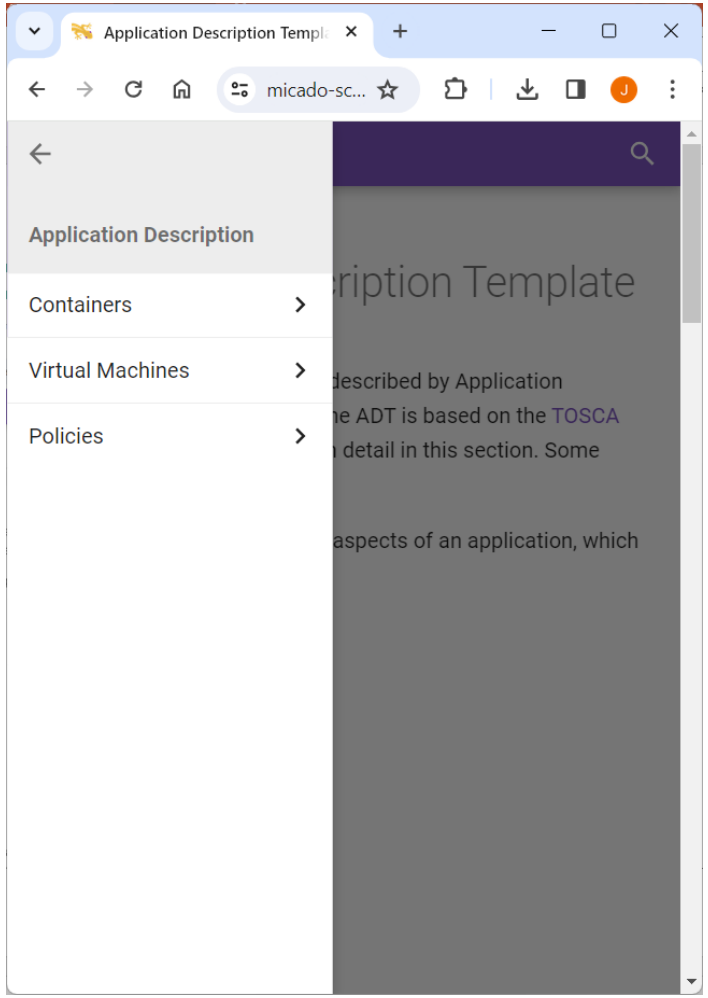

https://micado-scale.github.io

Summary

- Cloud-agnostic orchestration solution
- Pluggable architecture based on open-source components
- Standardised TOSCA-based application and policy description
- Automated application deployment in clouds
- Support for highly customisable scaling policies
- Support for large variety of clouds

CloudSigma

ORACLE

amazon web services

Google Cloud

Azure

EGI

CloudBroker

openstack.

https://micado-scale.github.io

Thank you for your interest!

https://micado-scale.github.io
presenter, designer, contributor: Jozsef Kovacs jozsef.kovacs@sztaki.hun-ren.hu
project leader: Tamas Kiss t.kiss@westminster.ac.uk
main developer: Jay Deslauriers j.deslauriers@westminster.ac.uk

Acknowledgement:
the work and slides presented were provided by the members of Centre of Parallel Computing at The University of Westminster