

Project COLA

Cloud Orchestration at the Level of Application

MiCADOscale is a highly customisable dynamic autoscaling framework that has been developed within Project COLA (Cloud Orchestration at the Level of Application). It particularly aims to increase the adoption of cloud computing services by SMEs and the public sector. Typical industry and public sector applications require resource scalability and efficient resource utilisation in order to serve a variable number of customers with dynamic resource demands, and to suitably optimise resource consumption and costs.

Project COLA is an Innovation Action funded by the European Commission as part of the Horizon2020 Programme. It started in January 2017 and lasts 33 months. The consortium includes 10 companies and 4 research organisations from 6 European countries, including the United Kingdom, Hungary, Sweden, Switzerland, Spain and Germany.



MiCADO

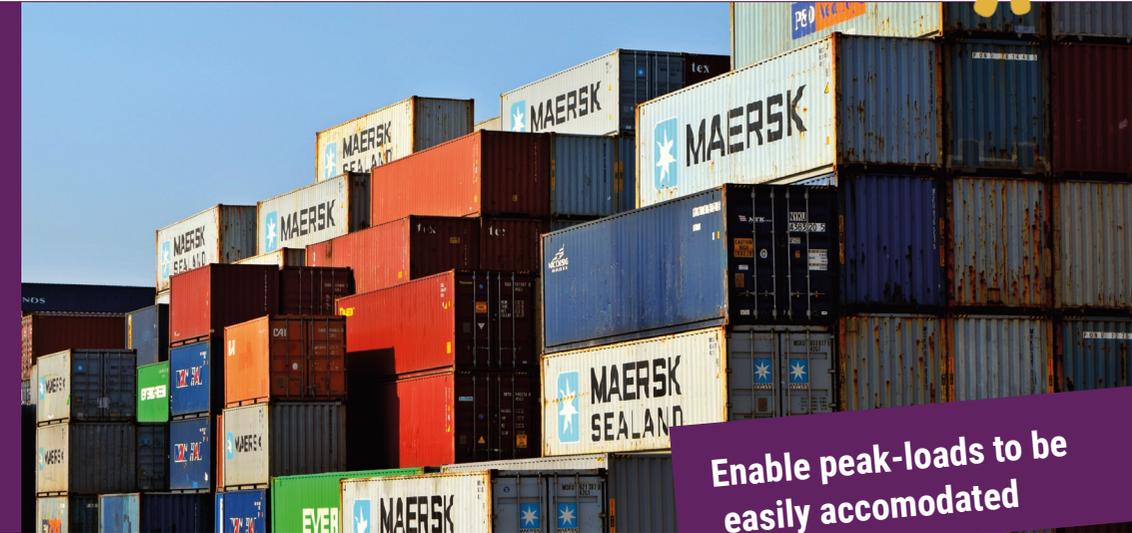
MiCADO

Auto-scaling Framework for Docker Containers
orchestrated by Kubernetes



MiCADO at a glance

- Open source modular framework
- Application level autoscaling (scaling both Docker containers and virtual machines)
- Multi Cloud Support, currently supported: AWS, CloudSigma, MS Azure, CloudBroker, OpenStack & OpenNebula
- Highly customable scaling policies
- Application level firewall by ZORP for enhanced security
- Multi-functional Dashboard (currently supported: Kubernetes dashboard, Grafana and Prometheus)
- professional support services for MiCADOscale by cloudSME



Enable peak-loads to be easily accommodated

Want consultation?



Prof. Tamas Kiss

Project Coordinator / Director of Research Centre for Parallel Computing, University of Westminster, London

t.kiss@westminster.ac.uk



Andreas Ocklenburg

CEO, cloudSME UG, Duisburg

andreas.ocklenburg@cloudsme.eu

+49 172 92 17 406

Contact:

cloudSME UG (haftungsbeschränkt)
TecTower
Bismarckstraße 142
47057 Duisburg
Germany
Phone: +49 (0)203 3639 9955
Email: info@cloudsme.eu

Web:

www.micado-scale.eu
www.project-cola.eu
www.twitter.com/MiCADO_EU
www.cloudsme.eu
www.facebook.com/cloudsme

11/09/2019



COLA
Cloud Orchestration
at the Level of Application

Responsible for this publication:



Competence Center for
Cloud-based HPC simulation

Project coordination:

UNIVERSITY OF
WESTMINSTER

Microservice-based Cloud Application-level Dynamic Orchestrator
Automated deployment and orchestration of
application cluster

MiCADO is a highly customisable multi-cloud orchestration and auto-scaling framework for Docker containers, orchestrated by Kubernetes and for Virtual Machines.

- * Open-source
- * Enhanced Security
- * Multi Cloud Support
- * Multi-functional Dashboard
- * Customisable scaling policies
- * Professional service for business



Project COLA is funded by the
Horizon 2020 framework programme



www.project-cola.eu

The sole responsibility of this publication lies with the authors. The European Union is not responsible for any use that may be made of the information contained therein. This Project is funded by the European Union under grant agreement no: 731574.



www.micado-scale.eu

Documentation, Tutorials, Demos

auto-scaling framework for Docker containers, orchestrated by Kubernetes

MiCADO can accommodate high peak loads while also optimising resource consumption by auto-scaling at the levels of virtual machines and Kubernetes deployments.

Industrial Technology Demonstrators

- ✓ **Deadline-based scalable Evacuation Service**
Complete evacuation simulation runs by a given deadline while optimising resource consumption.
- ✓ **Public Sector Social Media Analysis**
Process Big Data without worrying about peak loads and system crashes.
- ✓ **Server Resource & Peak-load Management**
Stop paying for underutilised resources, scale your resources horizontally



micado-scale.eu/demos



Read the Docs. Go to MiCADO's latest user documentation and discover our tutorials: stressng, cqueue, nginx and wordpress.

micado-scale.readthedocs.io

GitHub. Check out MiCADO's source code on GitHub! The MiCADO community is working on the next release that will introduce multi-user management, vertical scaling and High Availability. github.com/micado-scale

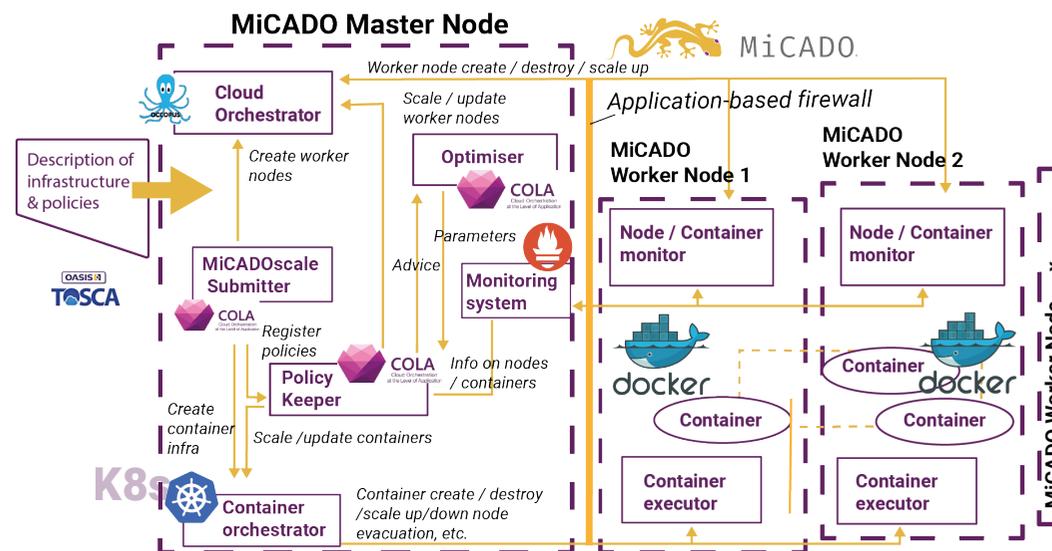


Point your smartphone's camera at the QR code and click on the notification URL that pops up.

Define your own scaling policies

automate deployment, scheduling and scaling of containerised applications

MiCADO core services must be deployed on a virtual machine (MiCADO Master Node) by an Ansible playbook*1. MiCADO Master is configured as the Kubernetes Master Node and has installed the Docker Engine, Occopus (to scale VMs), Prometheus (for monitoring), Policy Keeper (to perform decision on scaling) and Submitter (to provide submission endpoint) microservices to realise the autoscaling control loops. The application is deployed with an ADT, Application Description Template (yaml-file), and provides as IaC (Infrastructure as Code) all relevant information to roll out the application cluster framework.



During operation MiCADO workers (realised on new VMs) are instantiated on demand which deploy Prometheus, Node Exporter and CAdvisor as Kubernetes DaemonSets and the Docker engine through contextualisation. The newly instantiated MiCADO workers join the Kubernetes cluster managed by the MiCADO Master.

*1 creates the cloud resources to run the MiCADO Master Node, starts the MiCADO Master VM and deploys MiCADO to the VM

Become part of the MiCADOcommunity!

MiCADO is open-source: scan the code, visit our community on micado-scale.eu and discover demos, tutorials, publications and training sessions, developed by the University of Westminster and MTA SZTAKI. Feel free to test MiCADO, adjust it, make it better and share your changes with the community. MiCADO's code is hosted on GitHub and the documentation is provided on Read the Docs.



Add yourself to the **MiCADO Mailing List!**
micado-scale.eu/community/

MiCADOscale powered by cloudSME

Enjoy managed services, customised solutions directly, or by establishing connections with the right partners, e.g. infrastructure provider, technology provider or parallel computing scientists, within our affiliate network. Create seamlessly running services by the integration of MiCADOscale, delivering improved server capacity utilisation and greater flexibility allied to much reduced total operating costs.

Your European Development and Support Partner

Together, we aim to find the best solution for your individual case with the help of our affiliate network. Let's make an appointment for a chat to determine the future potentials of your applications and infrastructures.

