

DIGIT BRAIN

Living on the edge, sailing through the cloud

Orchestrating Applications in the Edge to Cloud Computing Continuum

Prof Dr Tamas Kiss

University of Westminster



t.kiss@westminster.ac.uk

 www.digitbrain.eu

 www.linkedin.com/groups/12439191

 www.twitter.com/digitbrain_eu

 www.facebook.com/DIGITbrainProject

DIGITbrain has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 952071



The Centre for Parallel Computing



Established in the 1990s

Focus on cloud/fog/edge computing

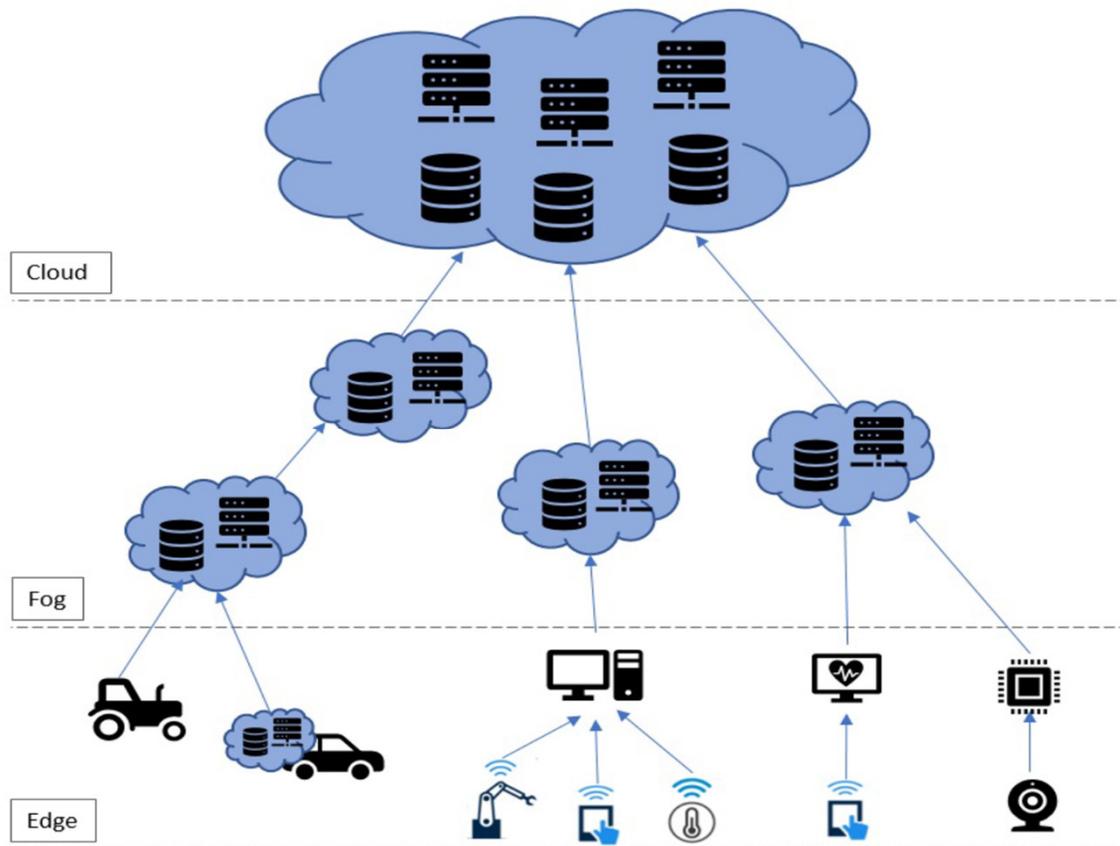
- Orchestration and autoscaling in the Cloud to Edge continuum
- Secure management of multi-cloud environments
- Cloud-native technologies – Docker, Kubernetes
- Standardised description of cloud-native applications - TOSCA
- Cloud/fog/edge based simulation environments for SMEs, research and the public sector
- Science/business gateways - user friendly interfaces for clouds and high performance computing

Main research outputs with impact in the past 6 years:

- MiCADO – application-level cloud to edge orchestrator - <https://micado-scale.eu/>
- emGORA workspace – a commercial digital marketplace for manufacturing simulation, data analytics and AI-based digital twins - <https://www.emgora.eu/>
- PITHIA e-Science Centre – Science Gateway for the space physics community - <https://esc.pithia.eu/>



The Cloud-to-Edge Compute Continuum



Application areas:

- Smart manufacturing/Industry 4.0
- Smart cities
- Precision agriculture
- Self-driving cars
- Augmented reality
- Etc.

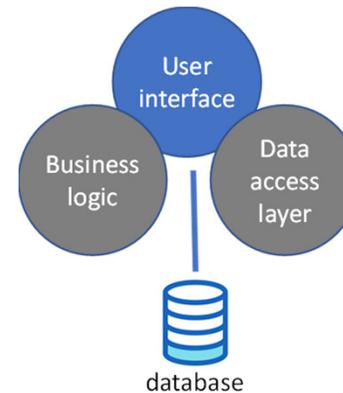


Monolithic vs Microservice architectures



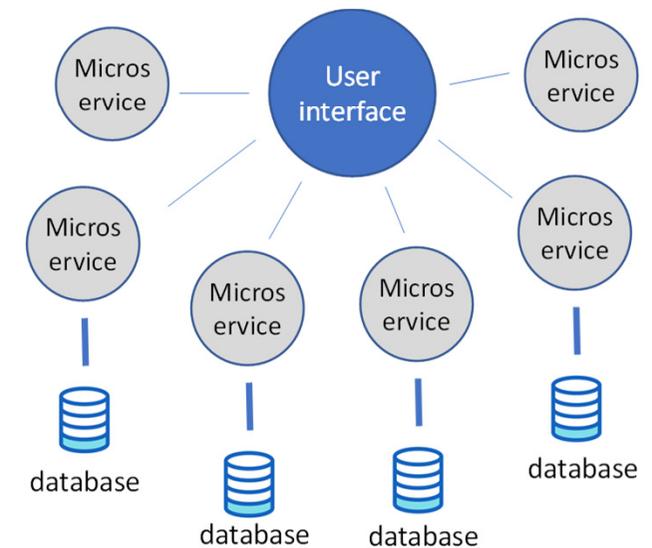
Monolithic

- ▮ All application components together
 - ▮ Single host, single DB, single server



Microservice

- ▮ Each application component individual
 - ▮ Containers make it possible!
 - ▮ Individual hosts, databases and servers
 - ▮ All working together



Cloud-native approach



- ❑ Cloud Native Computing Foundation - CNCF Cloud Native Definition v1.0:
<https://github.com/cncf/toc/blob/master/DEFINITION.md>
- ❑ Cloud native technologies empower organizations to build and run scalable applications in modern, dynamic environments such as public, private, and hybrid clouds.
- ❑ Containers, service meshes, microservices, immutable infrastructure, and declarative APIs exemplify this approach.



**CLOUD NATIVE
COMPUTING FOUNDATION**



Cloud-to-Edge orchestration



- Orchestration is the automated configuration, management, and coordination of computer systems, applications, and services.
- Cloud orchestration can be used to provision or deploy servers, assign storage capacity, create virtual machines, manage networking, deploy complex microservice architectures, auto-scale applications at run-time etc.
- Cloud-to-Edge orchestration extends this beyond the cloud layer; e.g. offload tasks from fog/edge servers to the cloud



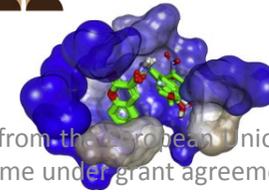
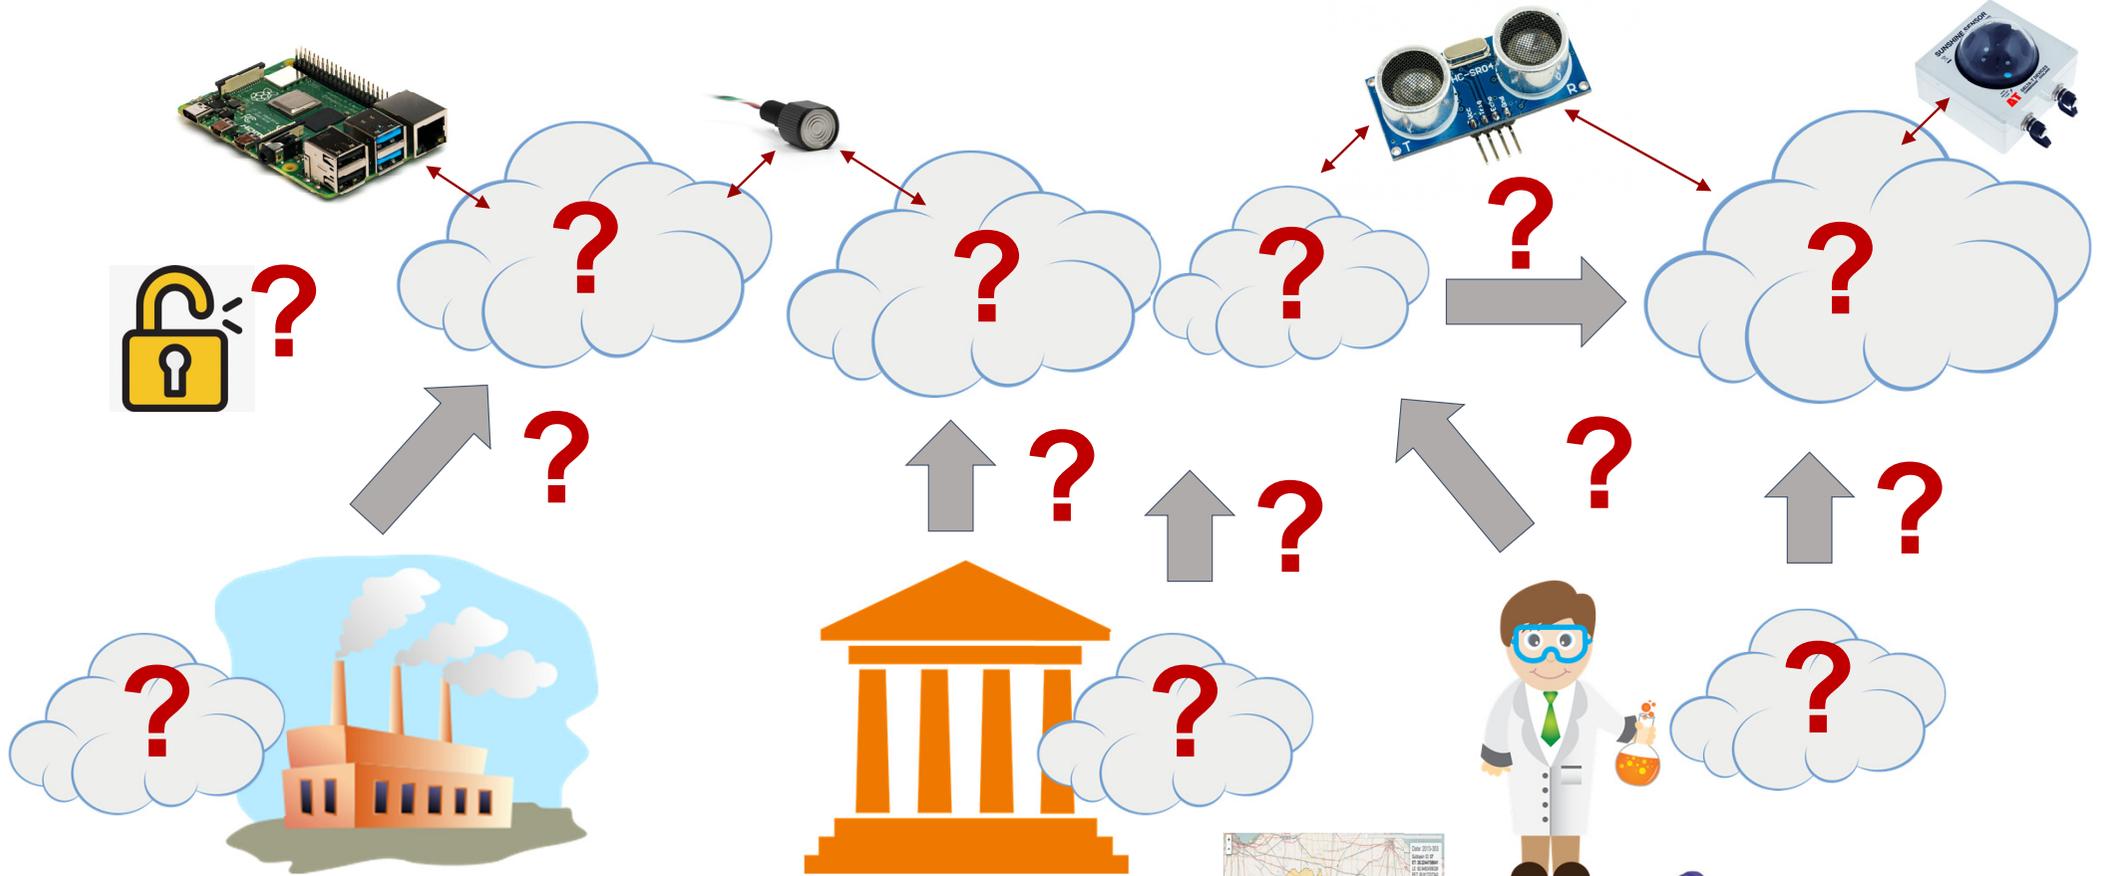
What challenges are we addressing?



- Cloud computing is gaining more and more attention in **various industrial sectors**
- **Private** clouds have significant benefits in terms of security and integrability into the enterprise environment but **hybrid** and multi-clouds are also widespread.
- New solutions are needed to support **Internet of Things (IoT)** and **Big Data** application areas -> cloud-fog-edge computing is a new challenge
- Growing demand for cloud orchestrators and brokering tools.



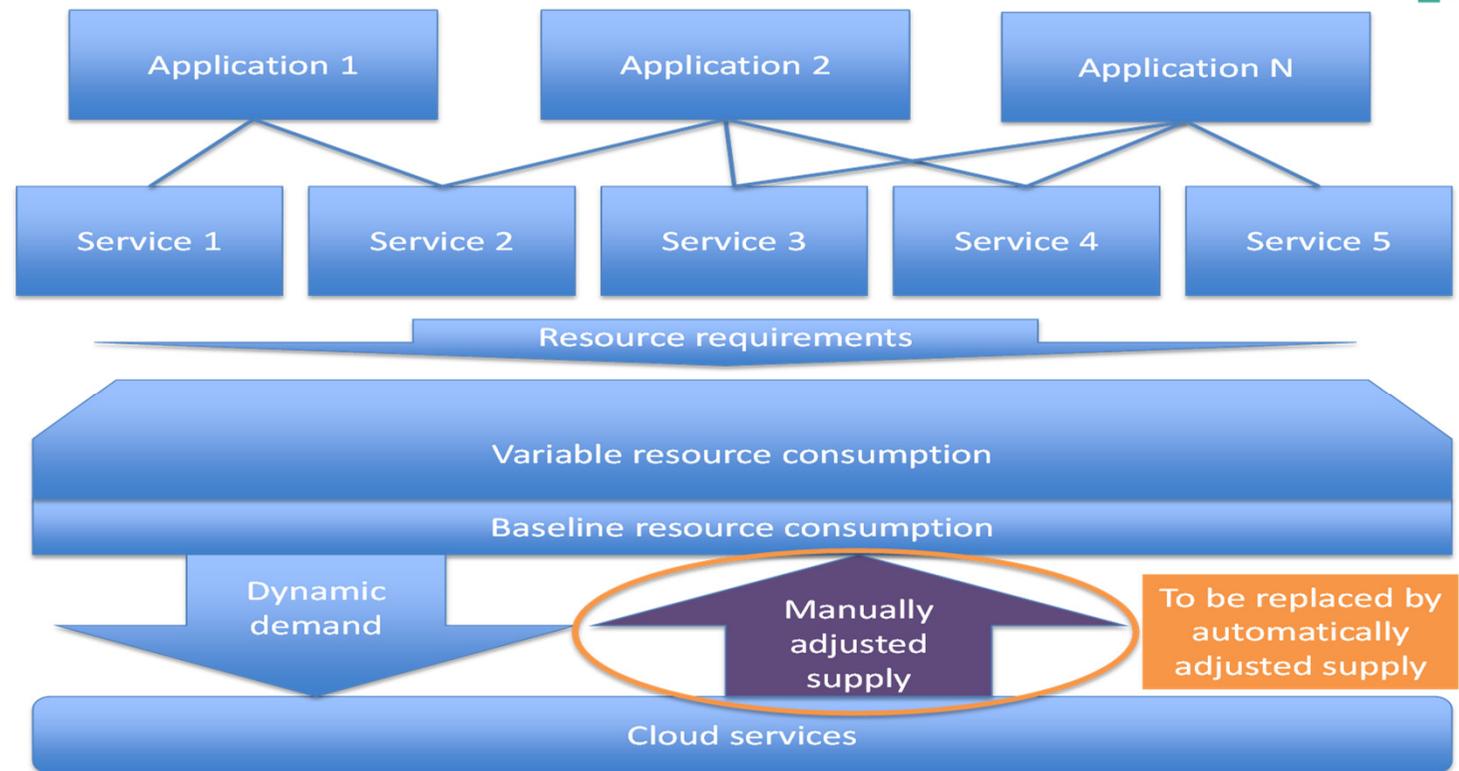
What challenges are we addressing?



The beginnings – application-level cloud orchestration



- ▮ deploying complex sets of application microservices in a cloud agnostic way
- ▮ supporting a wide range of scaling policies and intelligent scaling decisions
- ▮ supporting cloud application developers
- ▮ providing advanced policy-based security solutions



COLA – Cloud Orchestration at the Level of Application

<https://project-cola.eu/>

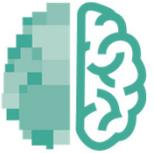


The COLA project

- EU H2020 project
- 1st January 2017 - 30th October 2019
- Funding: 4.2 million Euros
- 14 project partners from 6 European countries
- 10 companies and 4 academic/research institutions
- More information: <https://project-cola.eu>



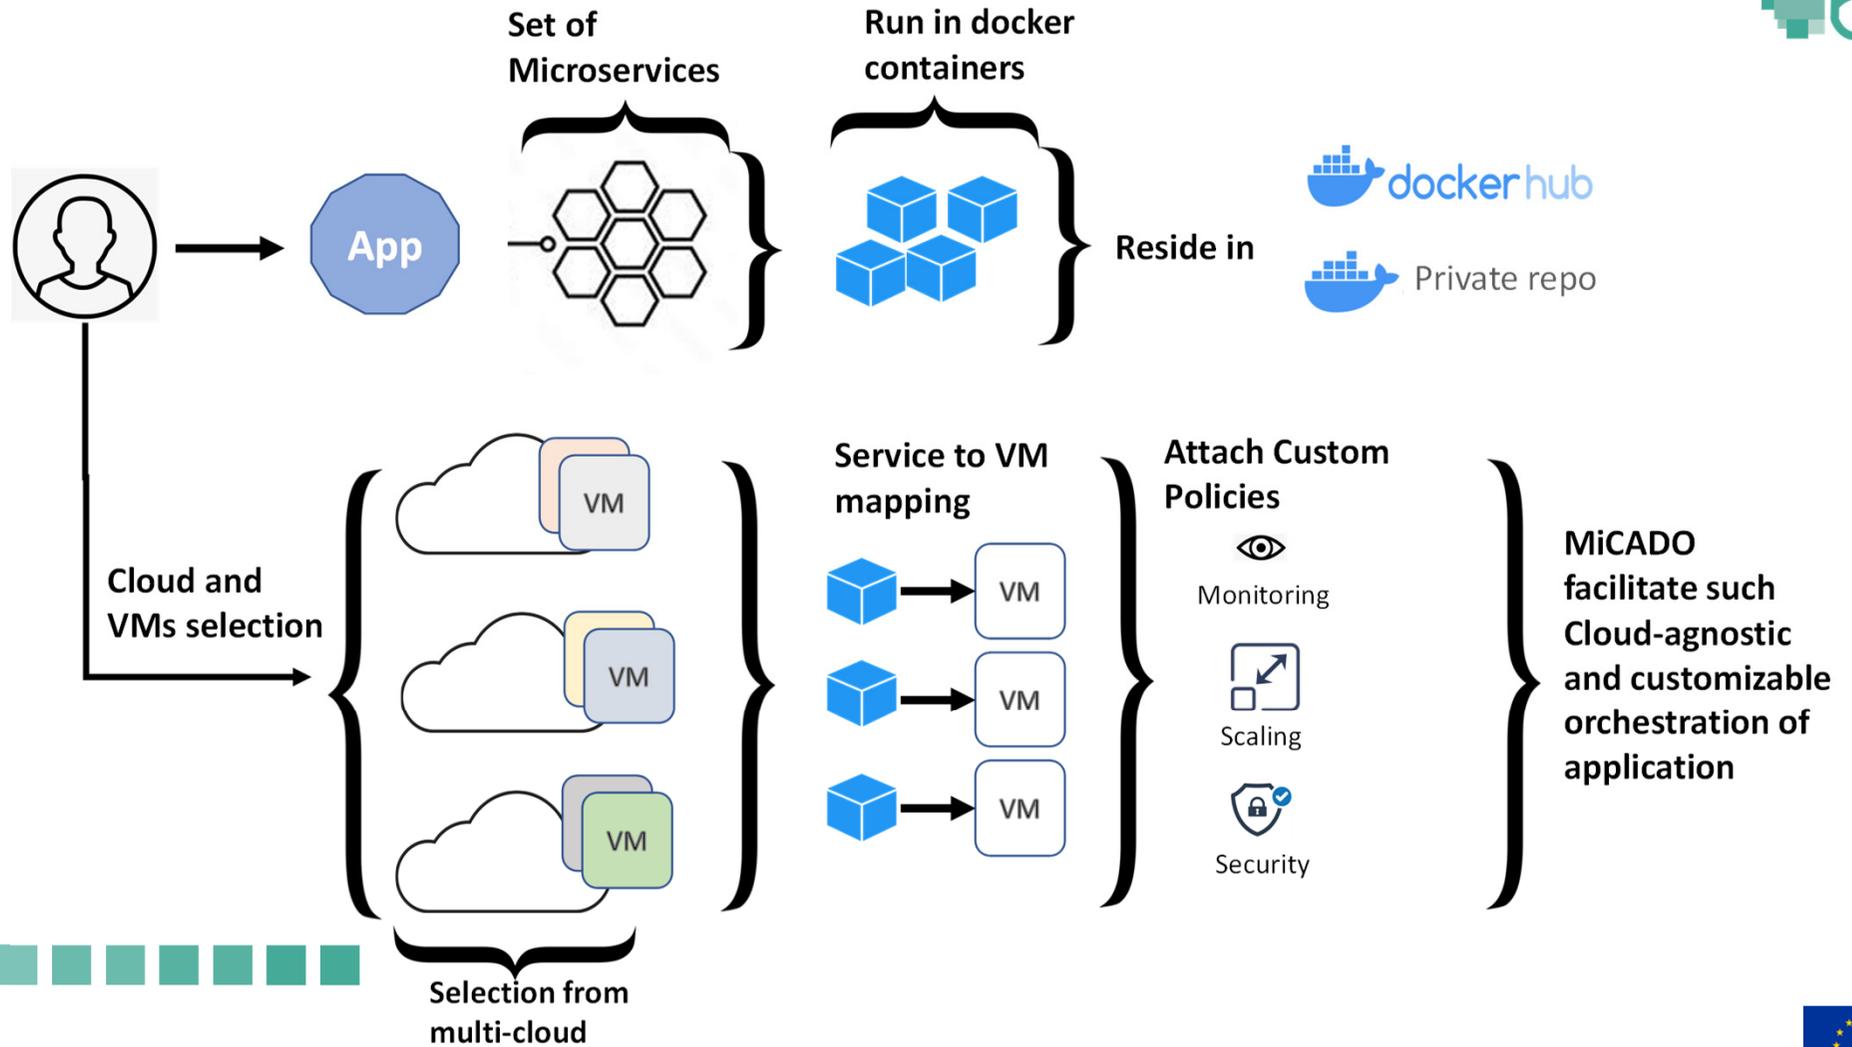
MiCADO — Microservices-based Cloud Application-level Dynamic Orchestrator



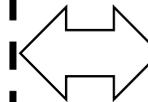
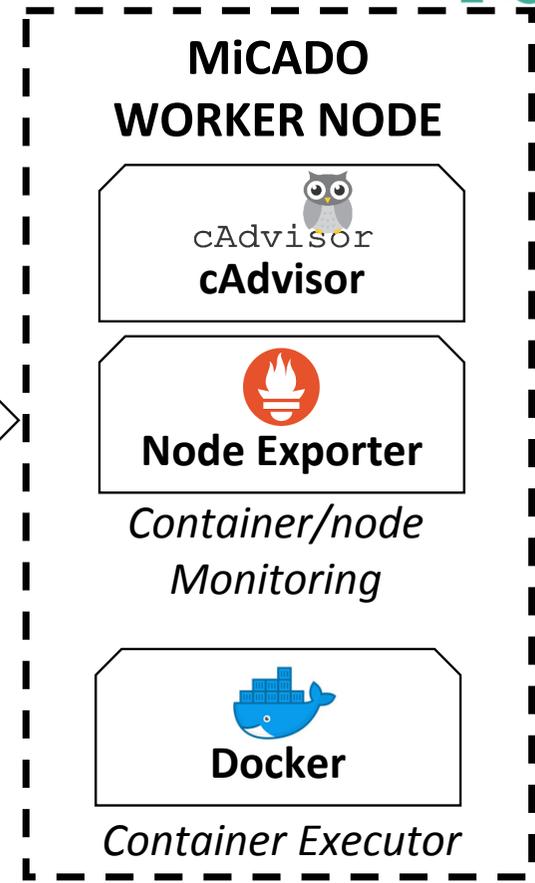
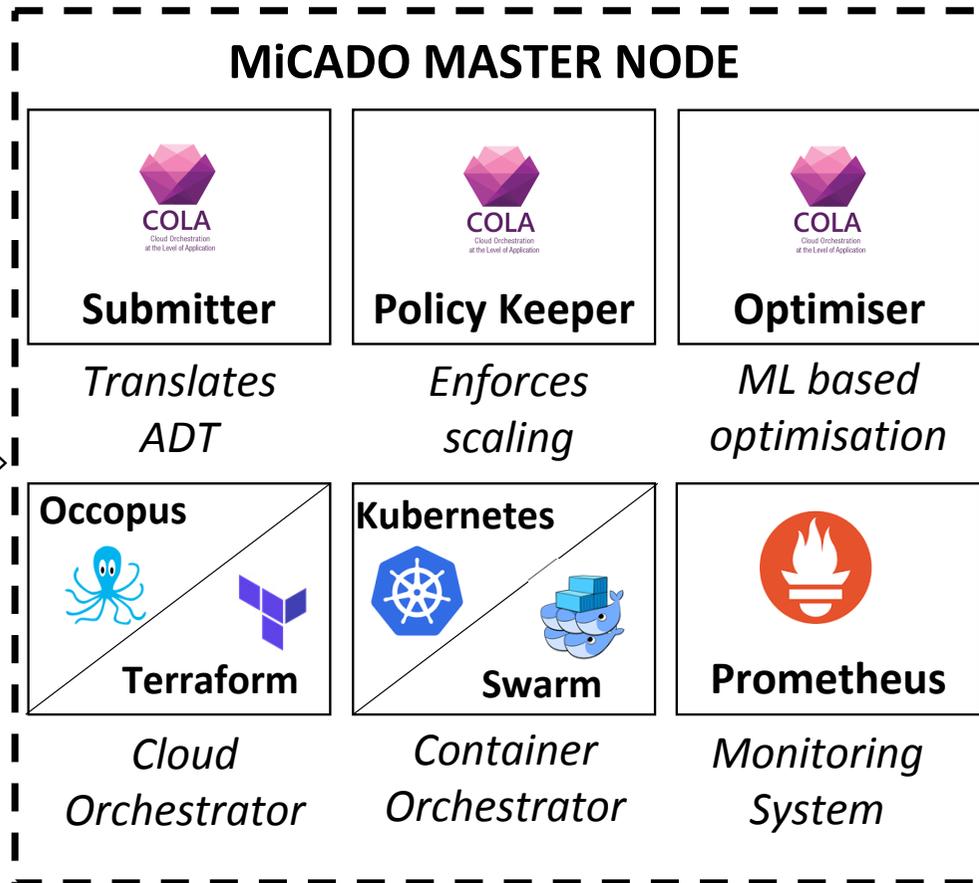
- ▮ Result of the H2020 COLA project
 - ▮ currently actively developed in the H2020 DIGITbrain project,
 - ▮ is/was used in ASCLEPIOS, PITHIA-NRF, CO-VERSATILE, Harpocrates and ARCAFF
- ▮ Automated application **deployment** based on TOSCA-based application description templates
- ▮ Automated **scaling** 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 - <https://micado-scale.eu/>



The overall concept



MiCADO — Microservices-based Cloud Application-level Dynamic Orchestrator



Application description in MiCADO



- ▮ Application Description Templates
- ▮ Written in Oasis Standard TOSCA
 - ▮ A cloud language in YAML
 - ▮ www.oasis-open.org/committees/tosca/
- ▮ Authored by developers
 - ▮ Understand application, metrics and scalable components
- ▮ Finalised by application operators
 - ▮ Provide Compute details for a Cloud Service Provider





Application description

- ▮ Three sections
 - ▮ Container Infrastructure
 - ▮ Containers, volumes, configurations
 - ▮ Cloud Infrastructure
 - ▮ Instance size, SSH keys, opened ports, VM image
 - ▮ Monitoring & Scaling Policy
 - ▮ Metric collection
 - ▮ Queries, alerts, thresholds, scaling logic

```
stressng:
  type: toasca.nodes.MiCADO.Container.Application.Docker.Deployment
  properties:
    image: lorel/docker-stress-ng
    args: ['--cpu', '0', '--cpu-method', 'pi', '-l', '20']
    resources:
      requests:
        cpu: "900m"

worker-node:
  type: toasca.nodes.MiCADO.CloudSigma.Compute.Occo.small
  properties:
    vnc_password: secret
    libdrive_id: ADD_LIBRARY_DRIVE_ID (e.g. 0d2dc532-39f5
    public_key_id: ADD_PUBLIC_KEY_ID (e.g. d7c0f1ee-40df-
    nics:
      - firewall_policy: ADD_FIREWALL_POLICY_ID (e.g. fd97e
        ip_v4_conf:
          conf: dhcp
```



Monitoring and autoscaling in MiCADO

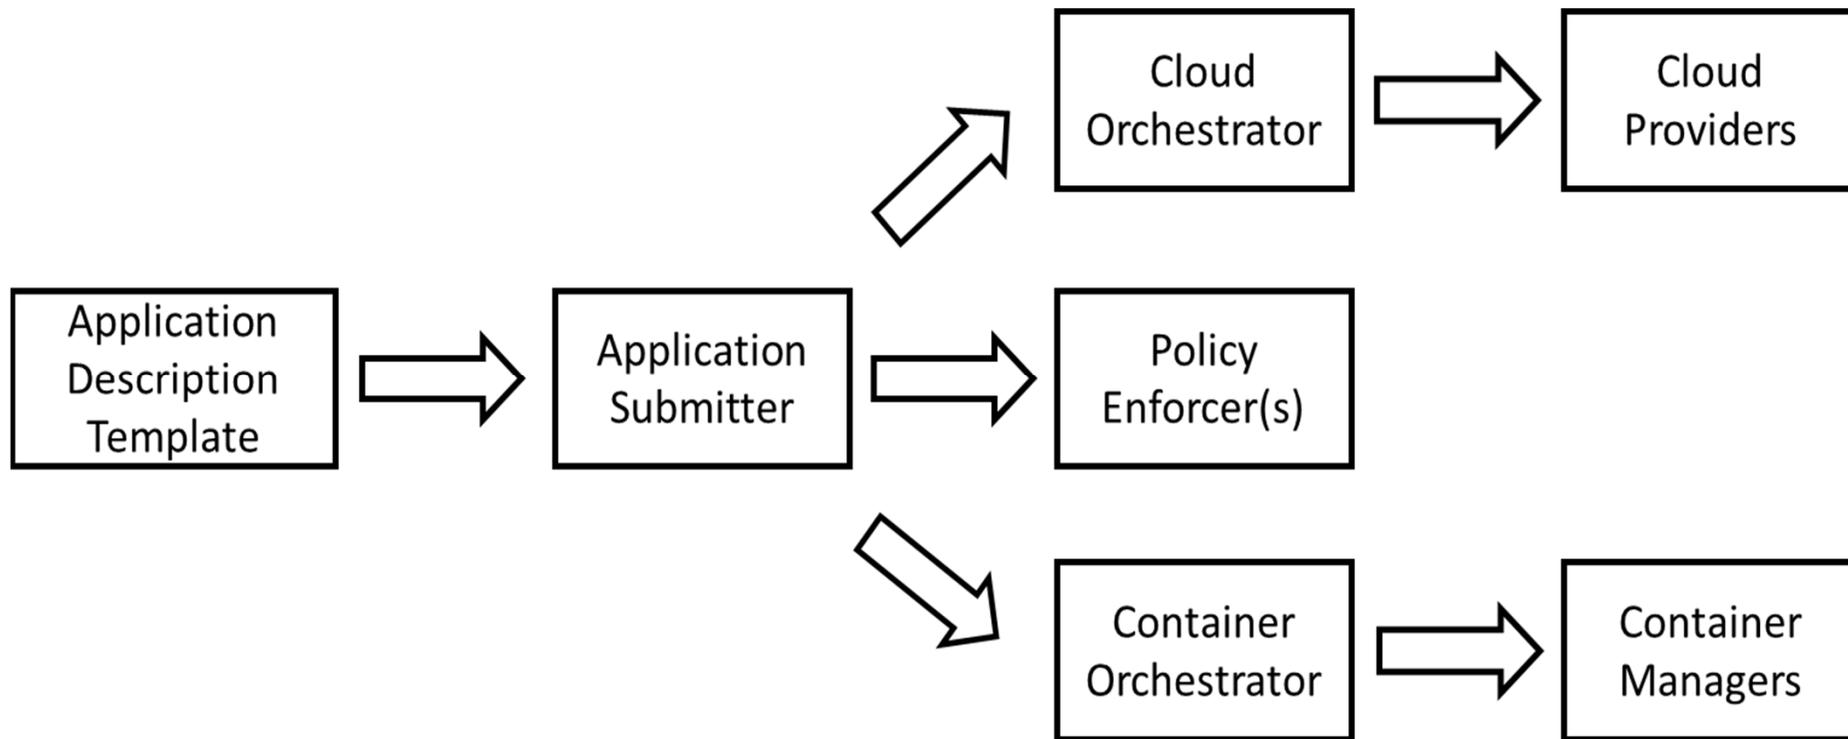


- Highly customizable monitoring subsystem
 - Monitored metrics are collected by dynamically attachable data collectors (Prometheus exporters)
- Highly customizable scaling logic
 - Scaling logic is fully programmable (using Python)
- Many different scaling policies are supported
 - Application types (job execution, web applications, ...)
 - Different metrics (cpu, network, number of jobs, ...)
 - Various strategies (load-based, deadline-based, event-based, ...)
- Scaling of containers and virtual machines are supported
 - Container-only and VM-only scaling
 - Scaling at both levels in parallel, independently or cooperatively
 - Possible to use predefined scaling policy or own-developed

```
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  
    .....
```

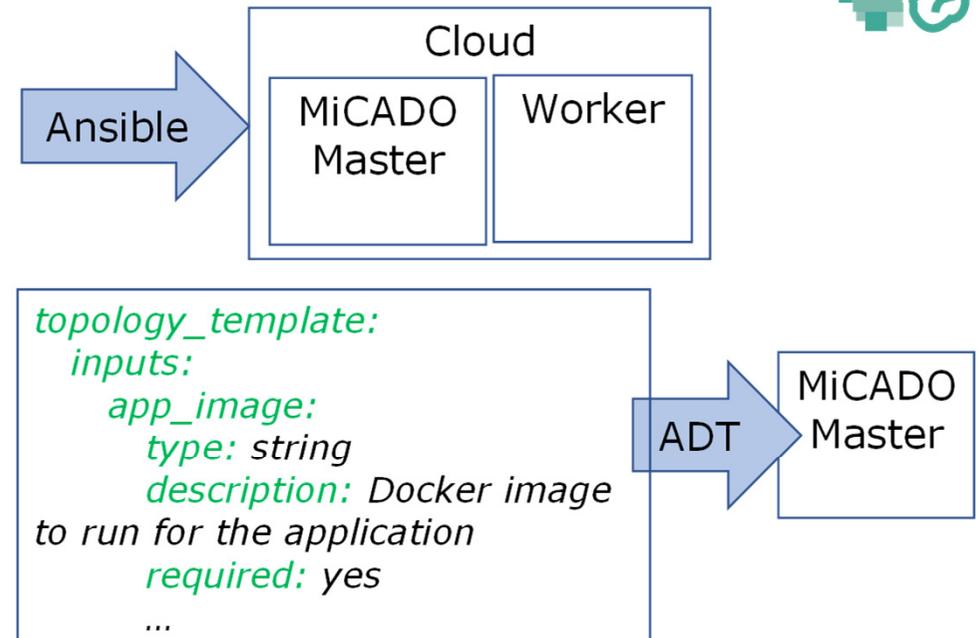


How does MiCADO work?

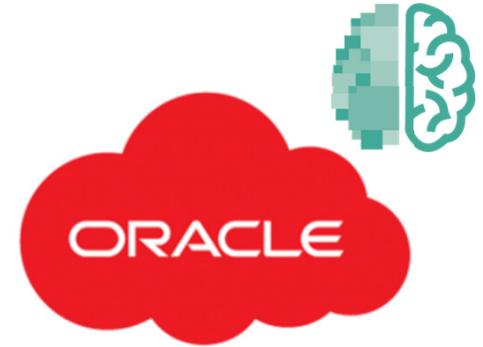


How to use MiCADO?

1. Deploy MiCADO by customizing Ansible configuration files
2. Describe your application (for e.g. virtual machine, scaling policy, etc.) by creating/ customizing TOSCA-based ADT file (Application Description Templates)
3. Submit your ADT file visa REST API call
4. Tracking MiCADO master and worker nodes in Dashboard during run-time



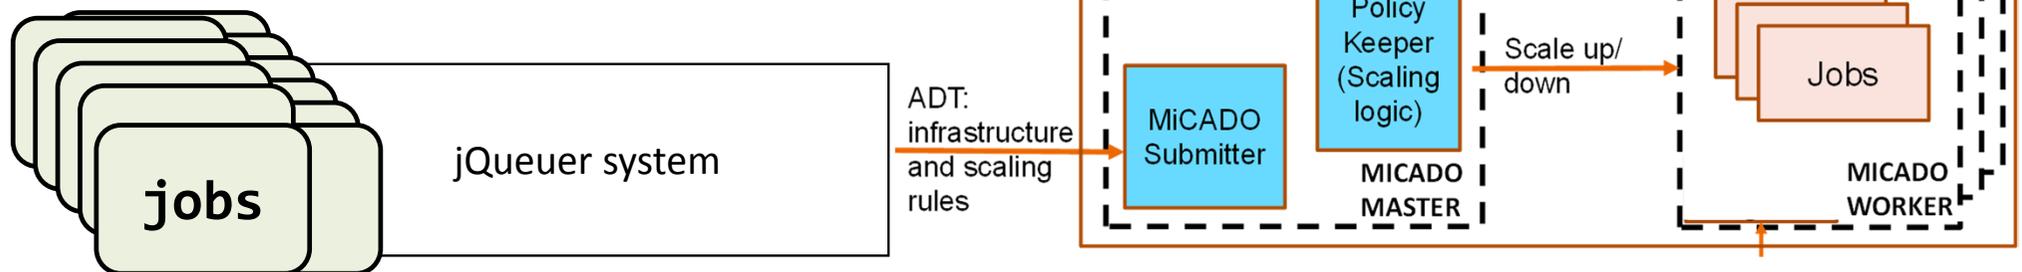
Supports a large variety of clouds





MiCADO and job execution

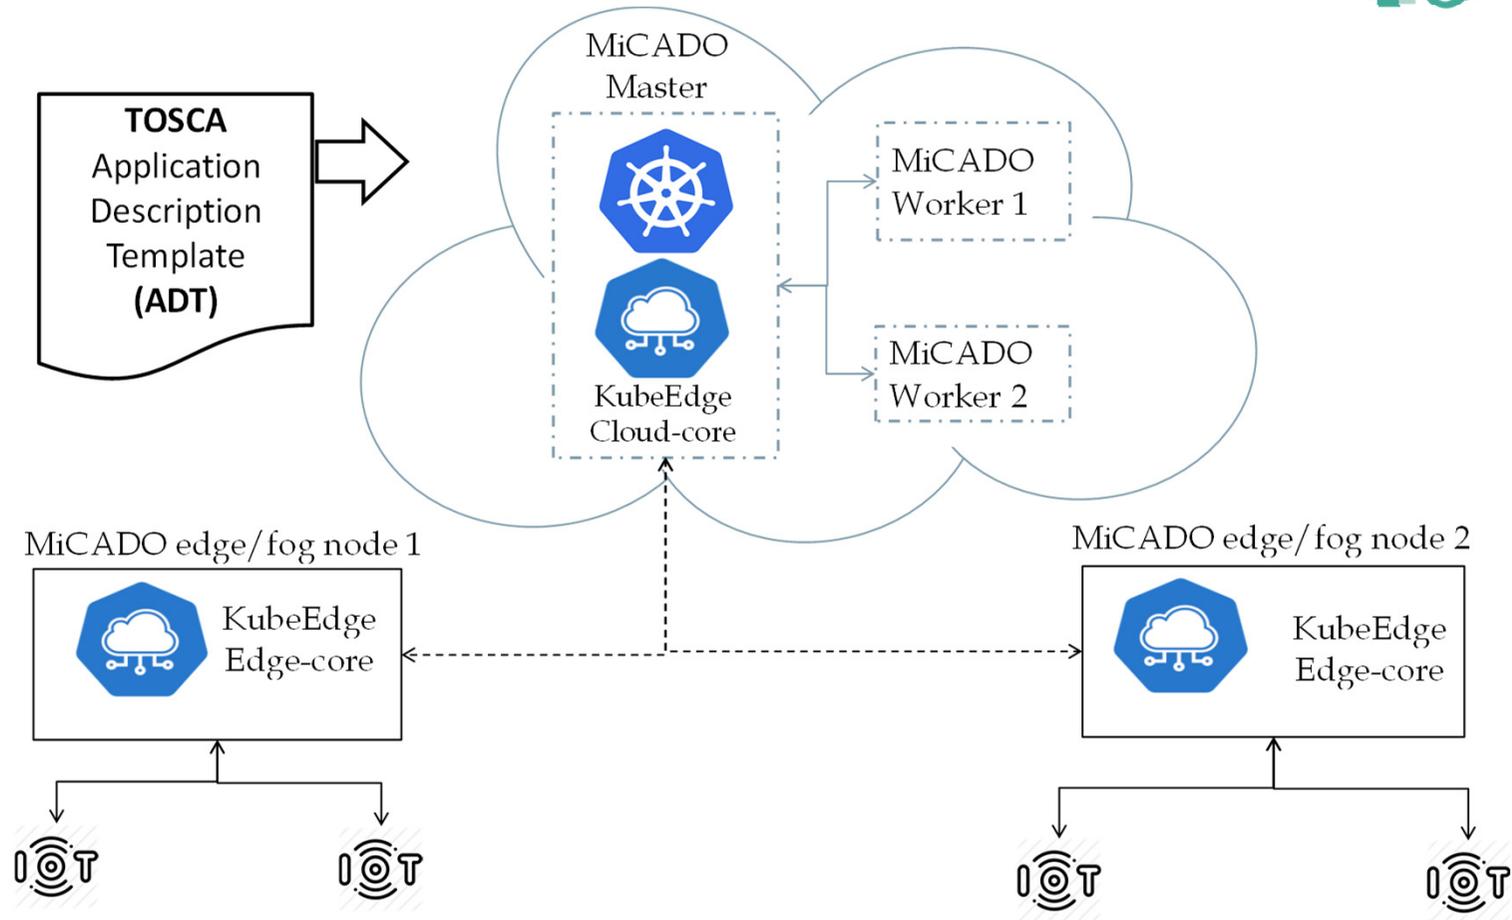
- Large number of jobs results in significant overall execution time
- Usually Restricted to complete all jobs by a deadline
 - Where to put the jobs?
 - How to distribute?
 - How to execute (in containers)
 - How to liaise with deadline?



How to extend orchestration to the Edge



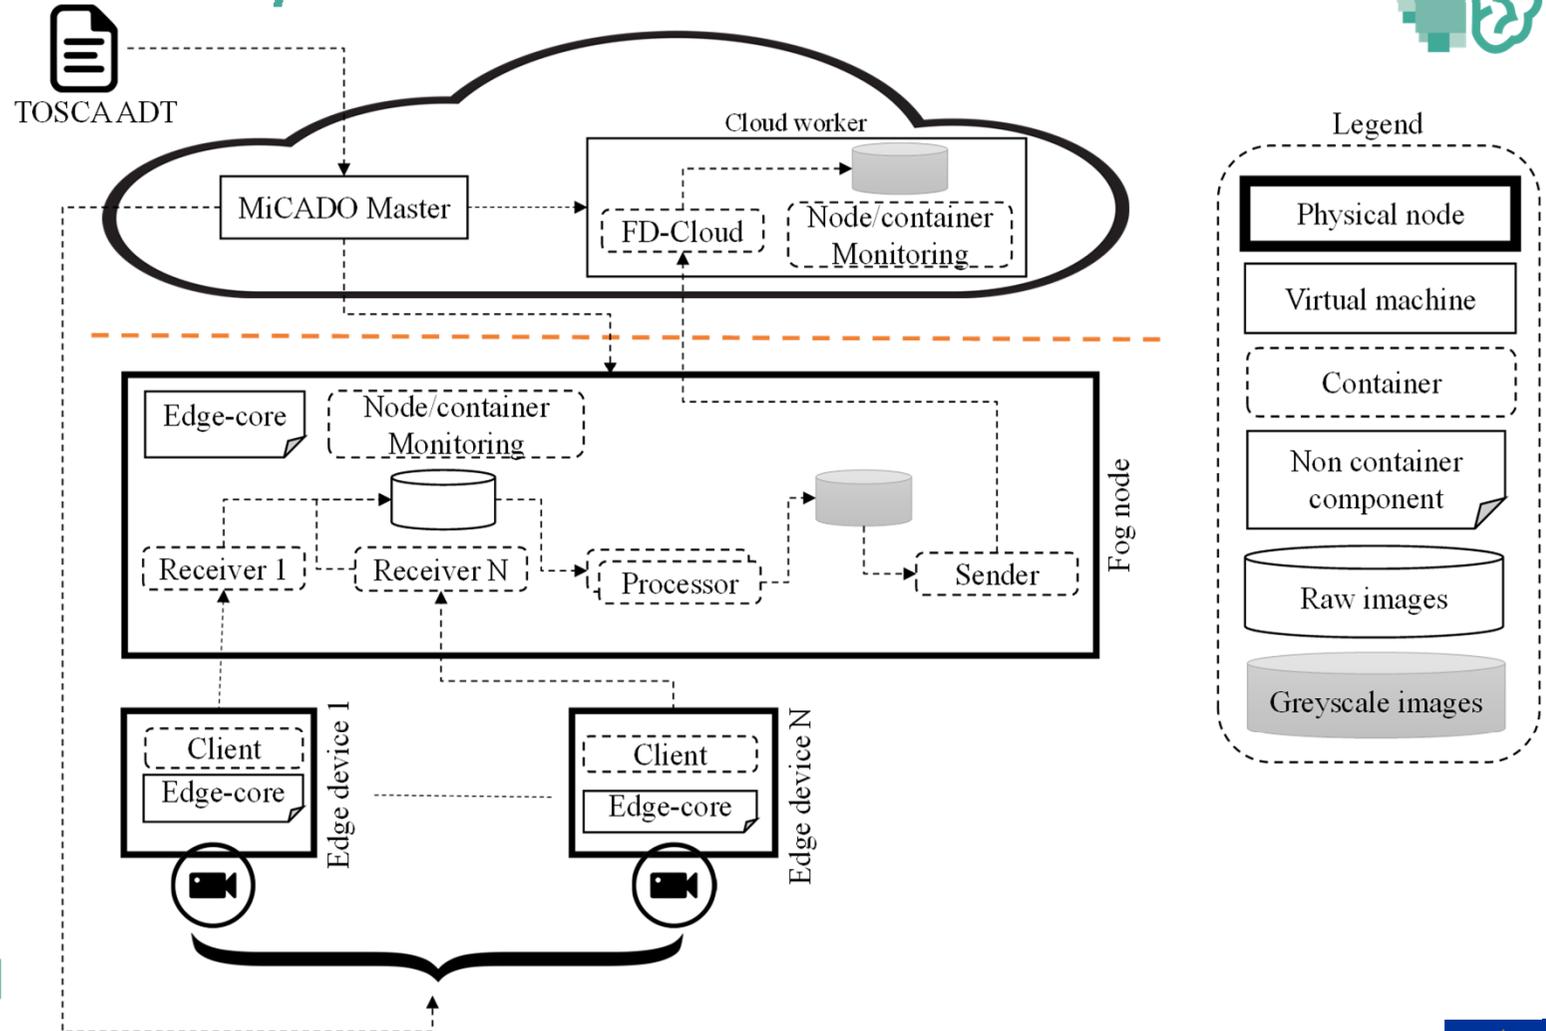
- ▮ Solution using KubeEdge
- ▮ Automated deployment of microservices extended to edge nodes
- ▮ Monitoring information collected from edge workers
- ▮ Scaling/reconfiguration policies extended towards edge



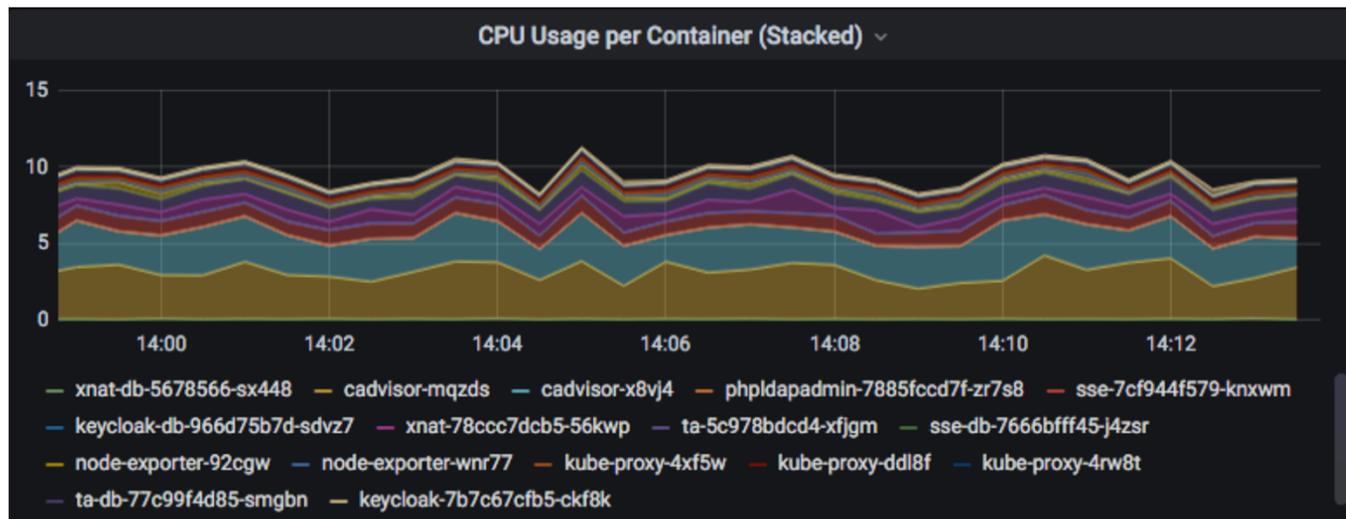
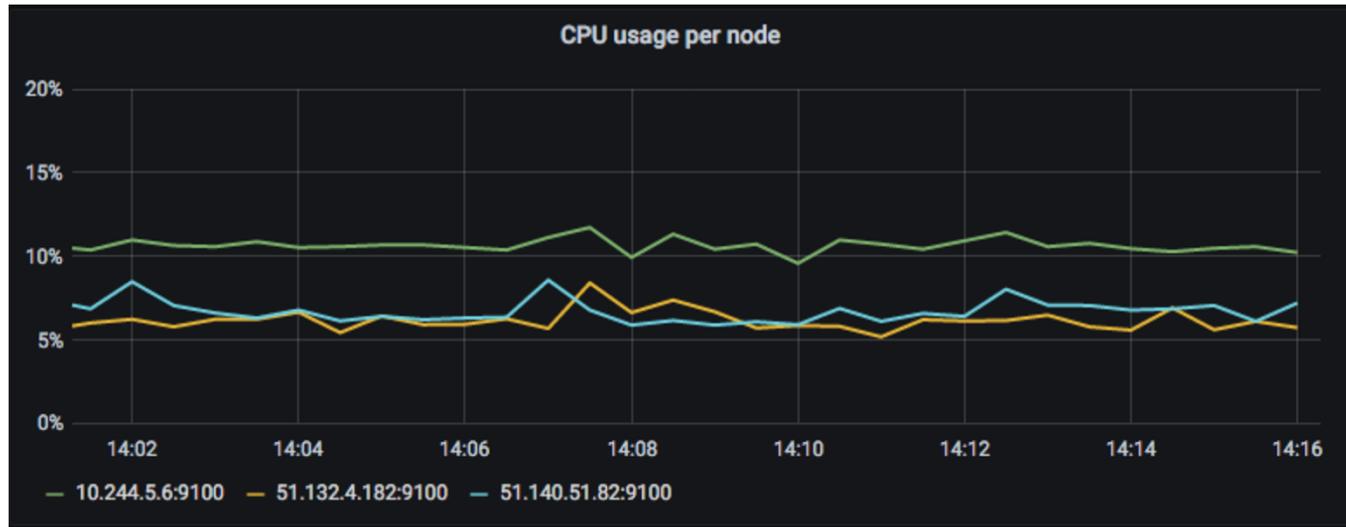
Face detection example



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



Face detection example MiCADO dashboard

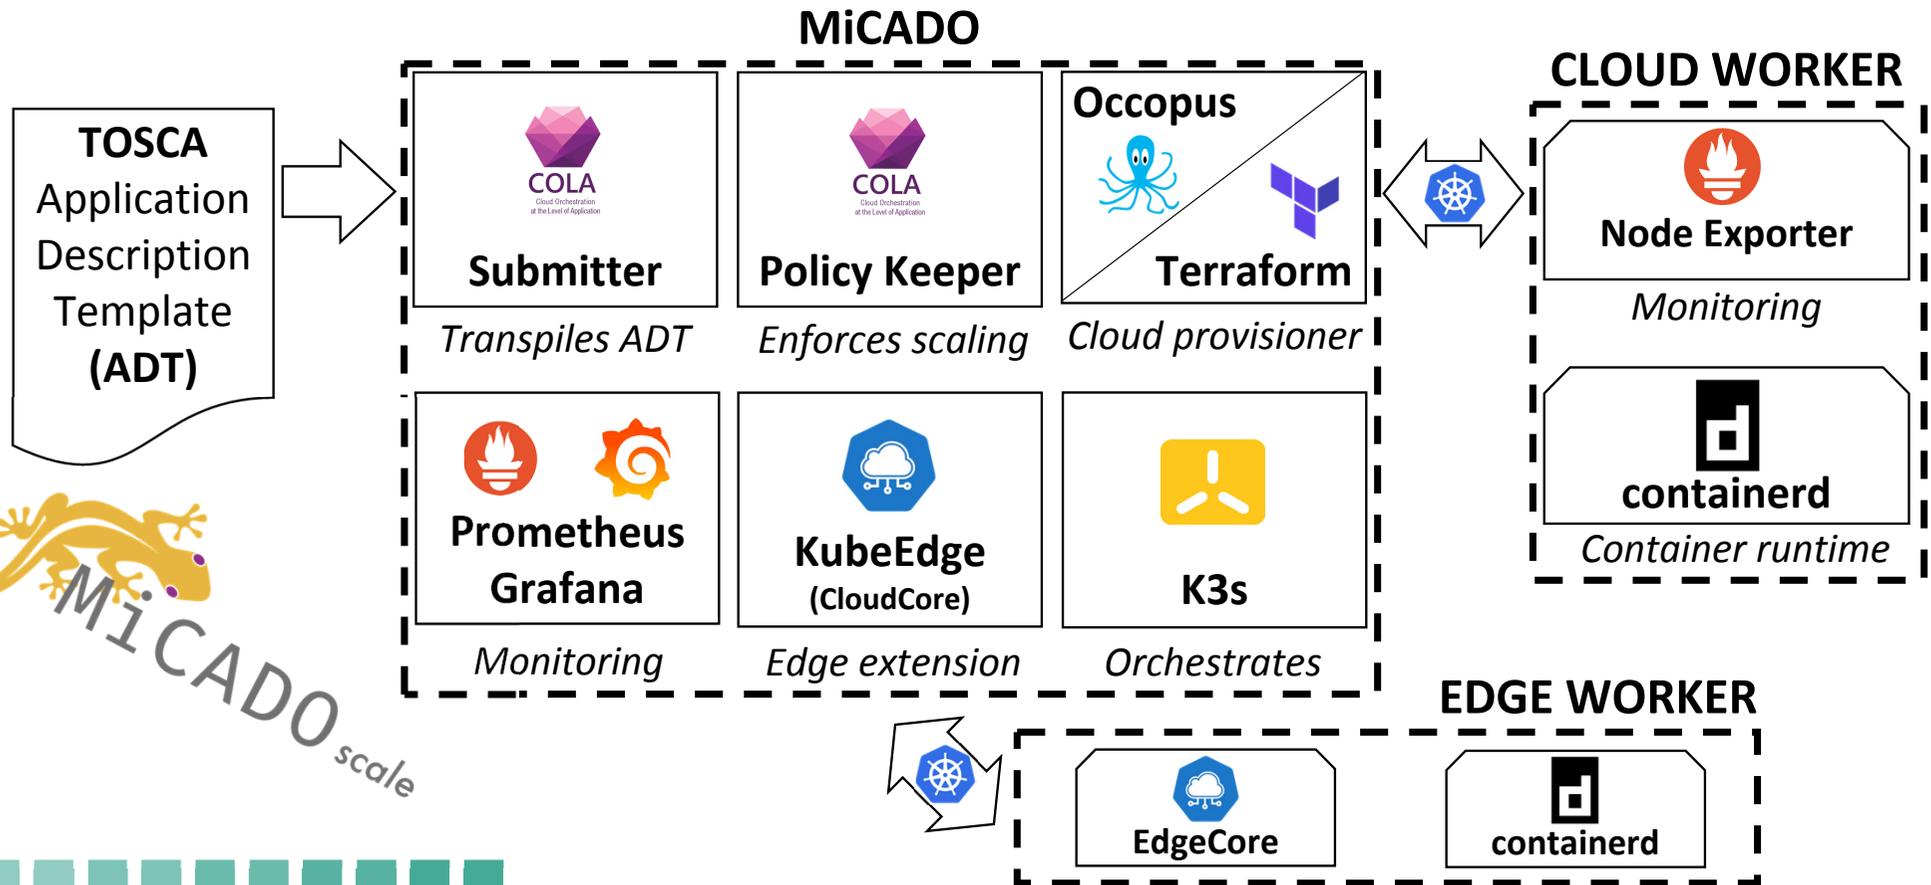


16/05/23

European Union's Horizon 2020
Agreement No 952071



Latest MiCADO architecture with Edge extension



MiCADO in the DIGITbrain project



Start date: 1st July 2020 - Duration: 42 months - Number of partners: 36

Core technical and administrative partners



Network of DIHs



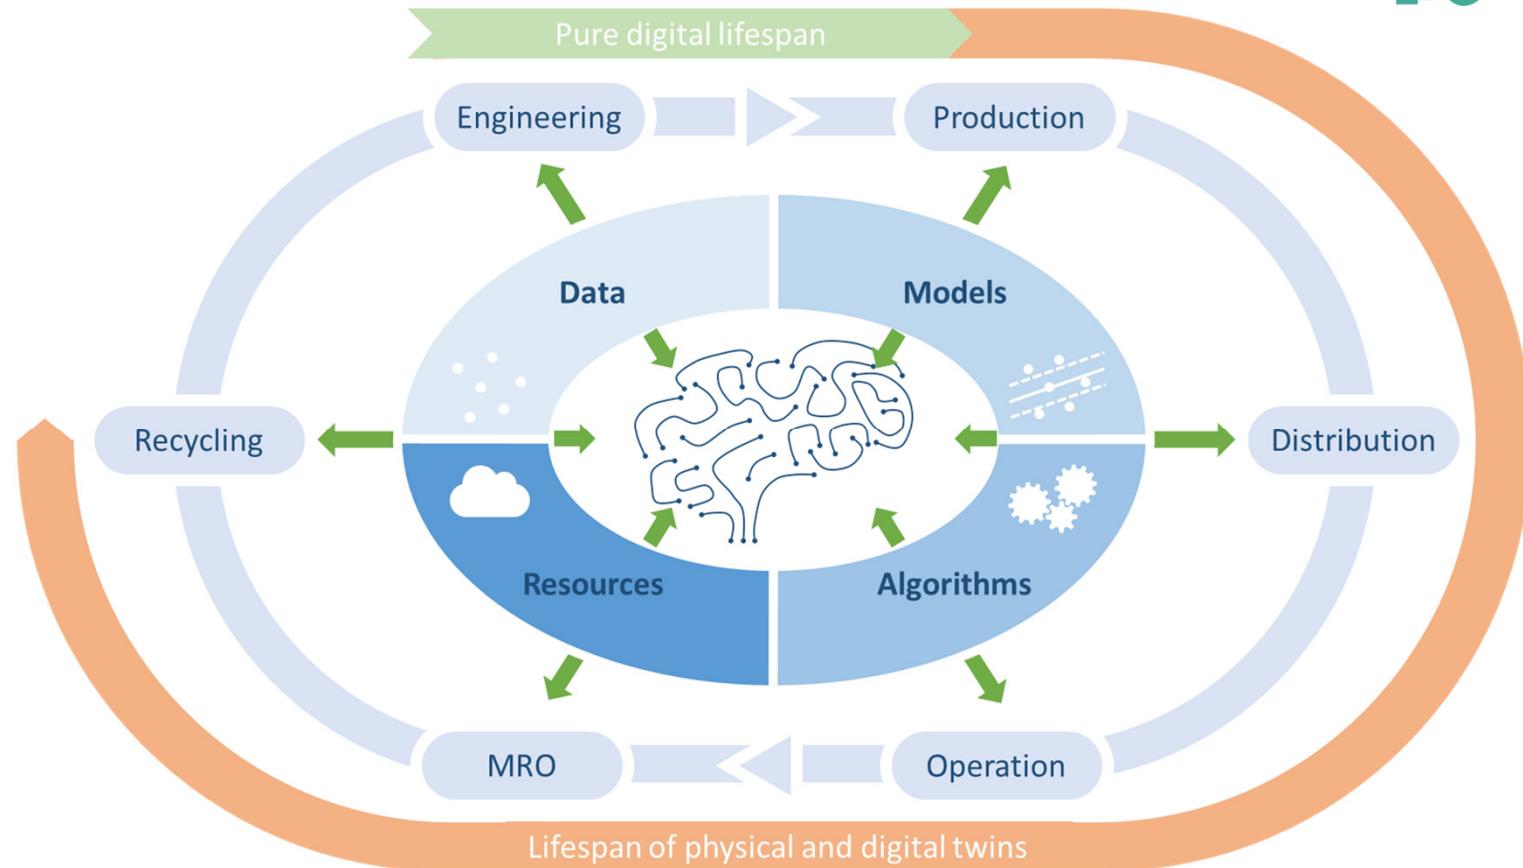
Experiment partners



MiCADO in the DIGITbrain project



- Compose Digital Twins from their basic building blocks (Data, Model, Algorithm)
- Execute Digital Twins in the Cloud-to-Edge continuum
- Collect, store and analyse data related to events in the Digital Product Brain



Key DIGITbrain concepts



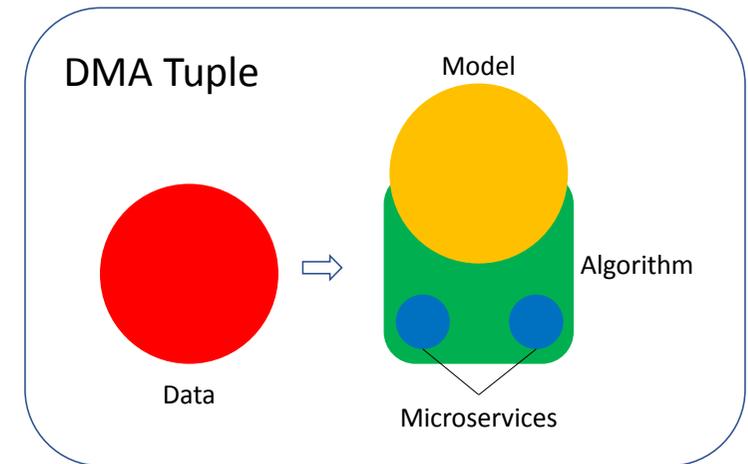
- ▮ **Industrial Products:** Mechatronic systems (or components) that are operated by manufacturing companies to support the production of other products (e.g. a manufacturing line).
- ▮ **Digital Twin:** Represented by static and dynamic models that can evaluate historical and actual data to assess the condition and to simulate the behaviour of the industrial product.
- ▮ **Digital Product Brain:** Guides the behaviour and performance of the industrial product by coalescing its physical and digital dimensions and by memorising the occurred (physical and digital) events over a significant part of its lifecycle.



Data – Model – Algorithm – Microservices



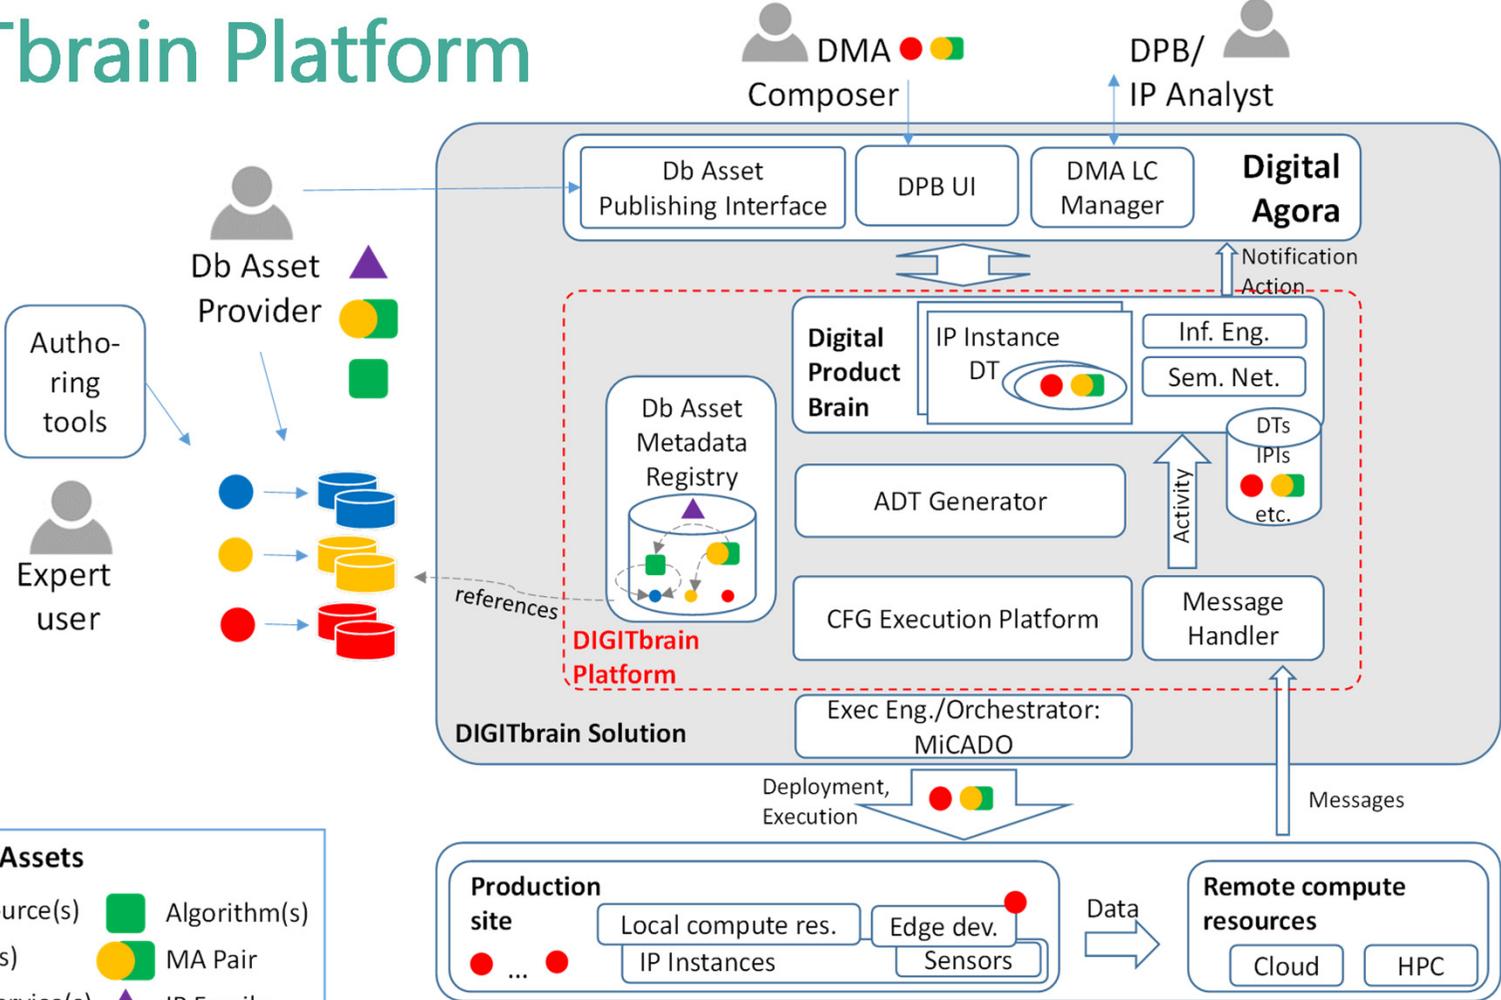
- ▮ **Data:** Refers to data sources in the form of files, streams or databases; data sources are considered to be external to the Db Platform (typically residing at the factory)
- ▮ **Model:** A Model is a description of a certain behaviour of/for an Industrial Product according to the given characteristics and operation conditions. A model is typically a file.
- ▮ **Microservice:** Encapsulates an executable procedure in a containerized form, e.g. in a Docker container.
- ▮ **Algorithm:** An Algorithm consists of one or more Microservice(s) and evaluate a model. Algorithms (its microservices) can be deployed into different resources (e.g. edge, cloud, HPC) to be executed, depending on the needs of the Model.



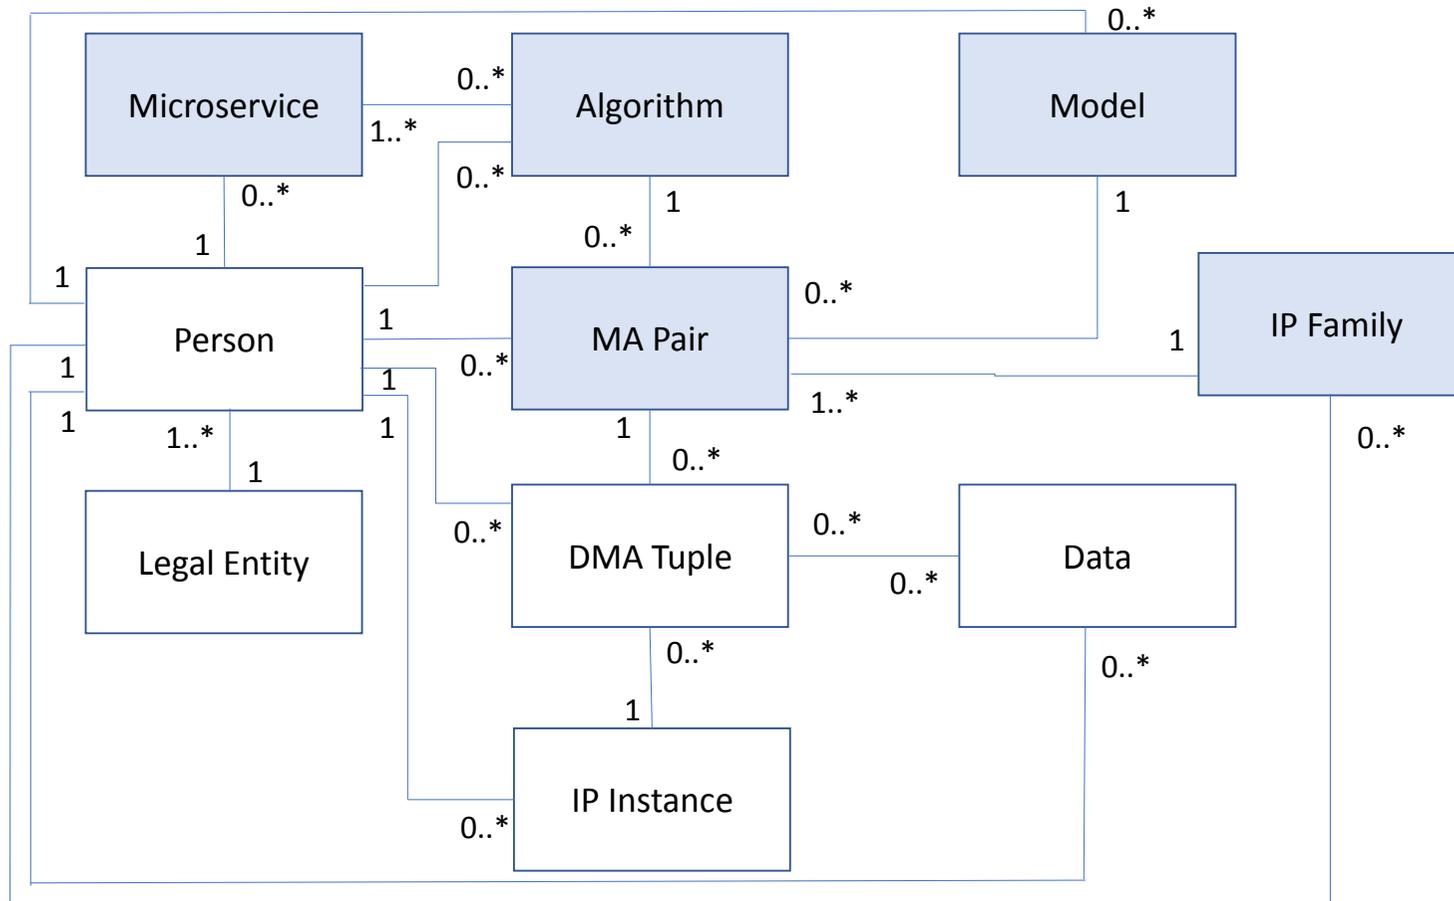
- ▮ Digital Twin: a DMA tuple



The DIGITbrain Platform



DIGITbrain Metadata



<https://digitbrain.github.io/>



DIGITbrain Metadata



Microservice

↑ Back to top

This page has been auto-generated based on an OpenAPI Specification
 Most users will prefer [this view](#) which may include examples and additional info

Name	Type	Details
name	string character varying	description: human readable short, yet descriptive name of the Microservice.
version	string character varying	description: version as defined by the user.
description	string character varying	description: human readable short description of the Microservice's capabilities.
classification_schema	string public.classification_schema	enum: [Simulation', 'ML', 'others'] description: fine-granular classification of the Microservice
type	string ARRAY	description: detailed type of the microservice, list of keywords
deployment_format	string public.deployment_format	enum: [docker-compose', 'kubernetes-manifest'] description: identifier of the deployment environment required to deploy the Microservice's container
deployment_data	string json	description: JSON of docker-compose or kubernetes manifest required to run the container
configuration_data	string ARRAY	description: List of objects specifying configuration file(s) content required by the service

Algorithm

This page has been auto-generated based on an OpenAPI Specification
 Most users will prefer [this view](#) which may include examples and additional info

Name	Type	Details
name	string character varying	description: a human-readable name to ease identification and discoverability for human users
description	string character varying	description: a short, human-readable description of the Algorithm to aid a human user in analysing the Algorithm's capabilities and its applicability to a certain problem
classification_schema	string public.classification_schema	enum: [Simulation', 'ML', 'others'] description: the classification of the Algorithm, to describe the specialization area
type	string ARRAY	description: a detailed list of attributes to describe the Algorithm's field of application
version	string character varying	description: the version, as defined by the provider
list_of_microservices	string ARRAY	description: a list of Microservice Asset IDs, which are contained in the algorithm
deployment_mapping	string json	description: a mapping specifying which microservice should run on which host. By default each microservice is assigned a respective host, but this behaviour is not always ideal (eg. when two or more Microservices may need to run on the same host)



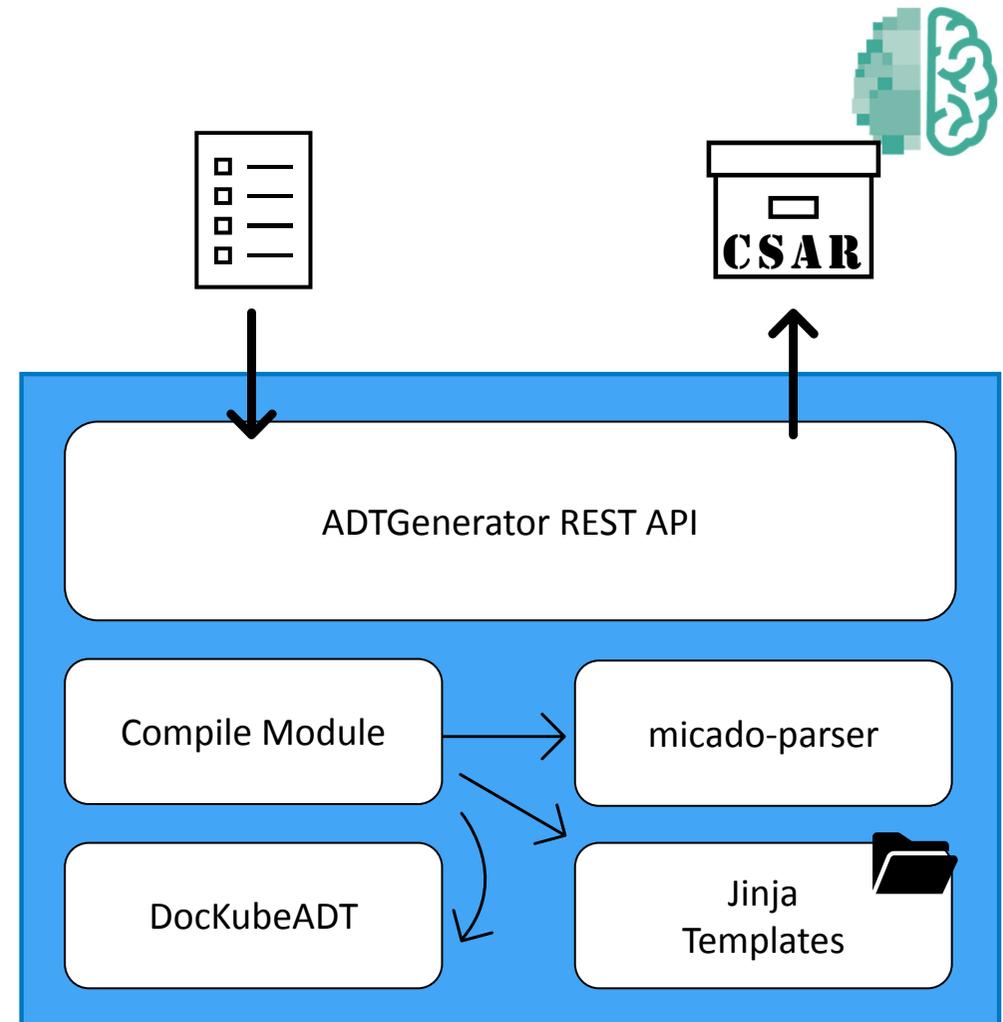
The ADT Generator

▮ Aim:

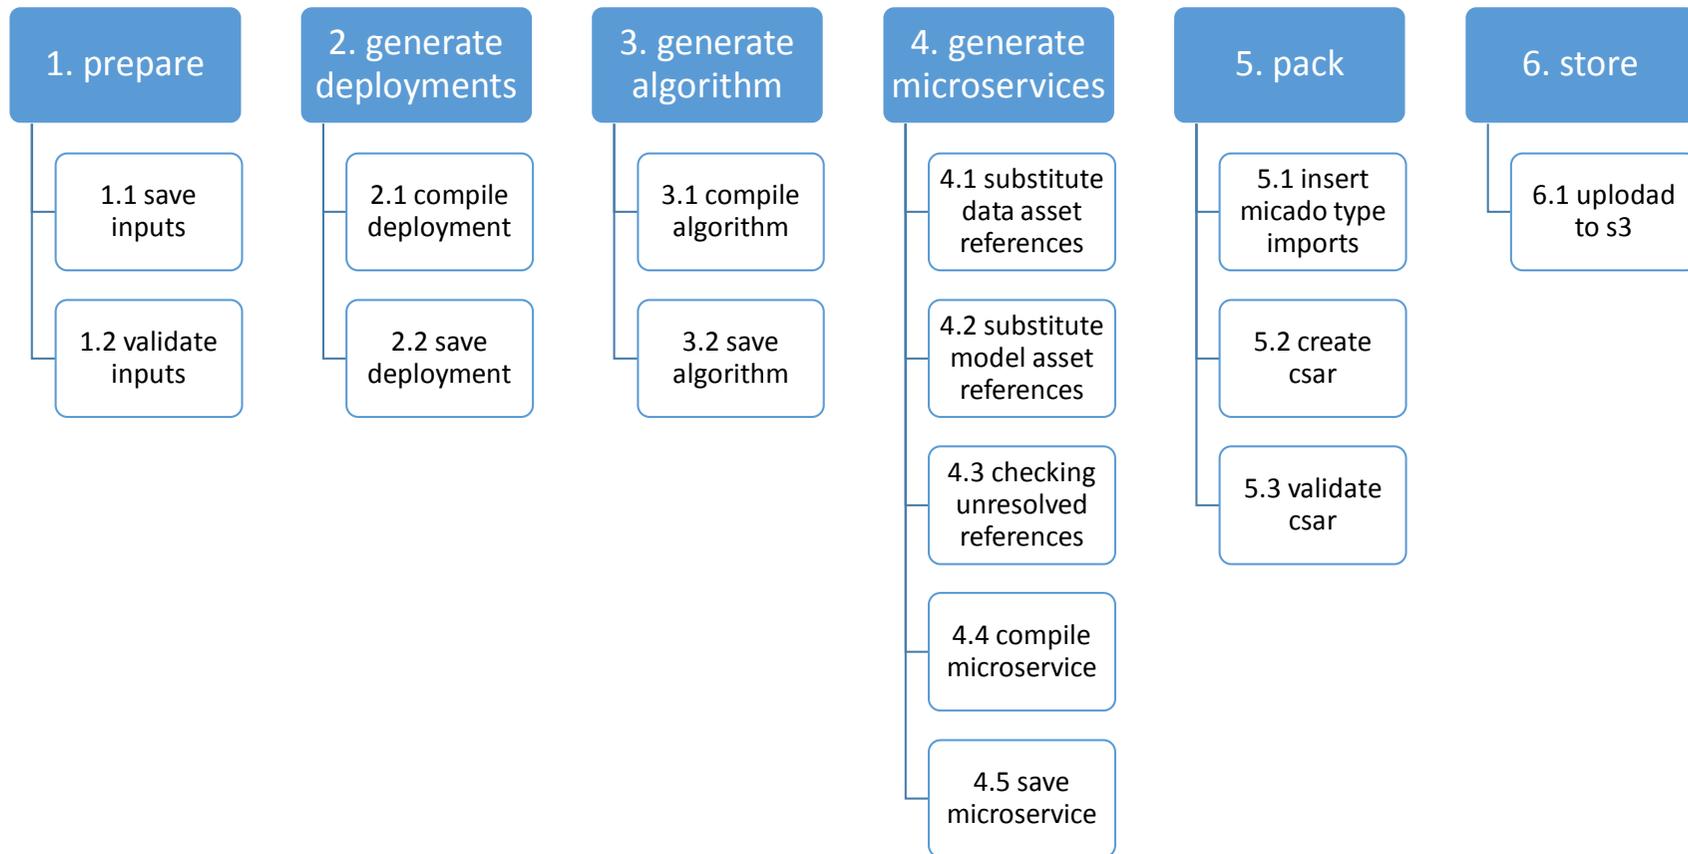
- ▮ Translate metadata key/value pairs into a TOSCA application descriptor
- ▮ Compose various metadata (Ms, D, M, A, MA, DMA) into a single descriptor file

▮ Challenge:

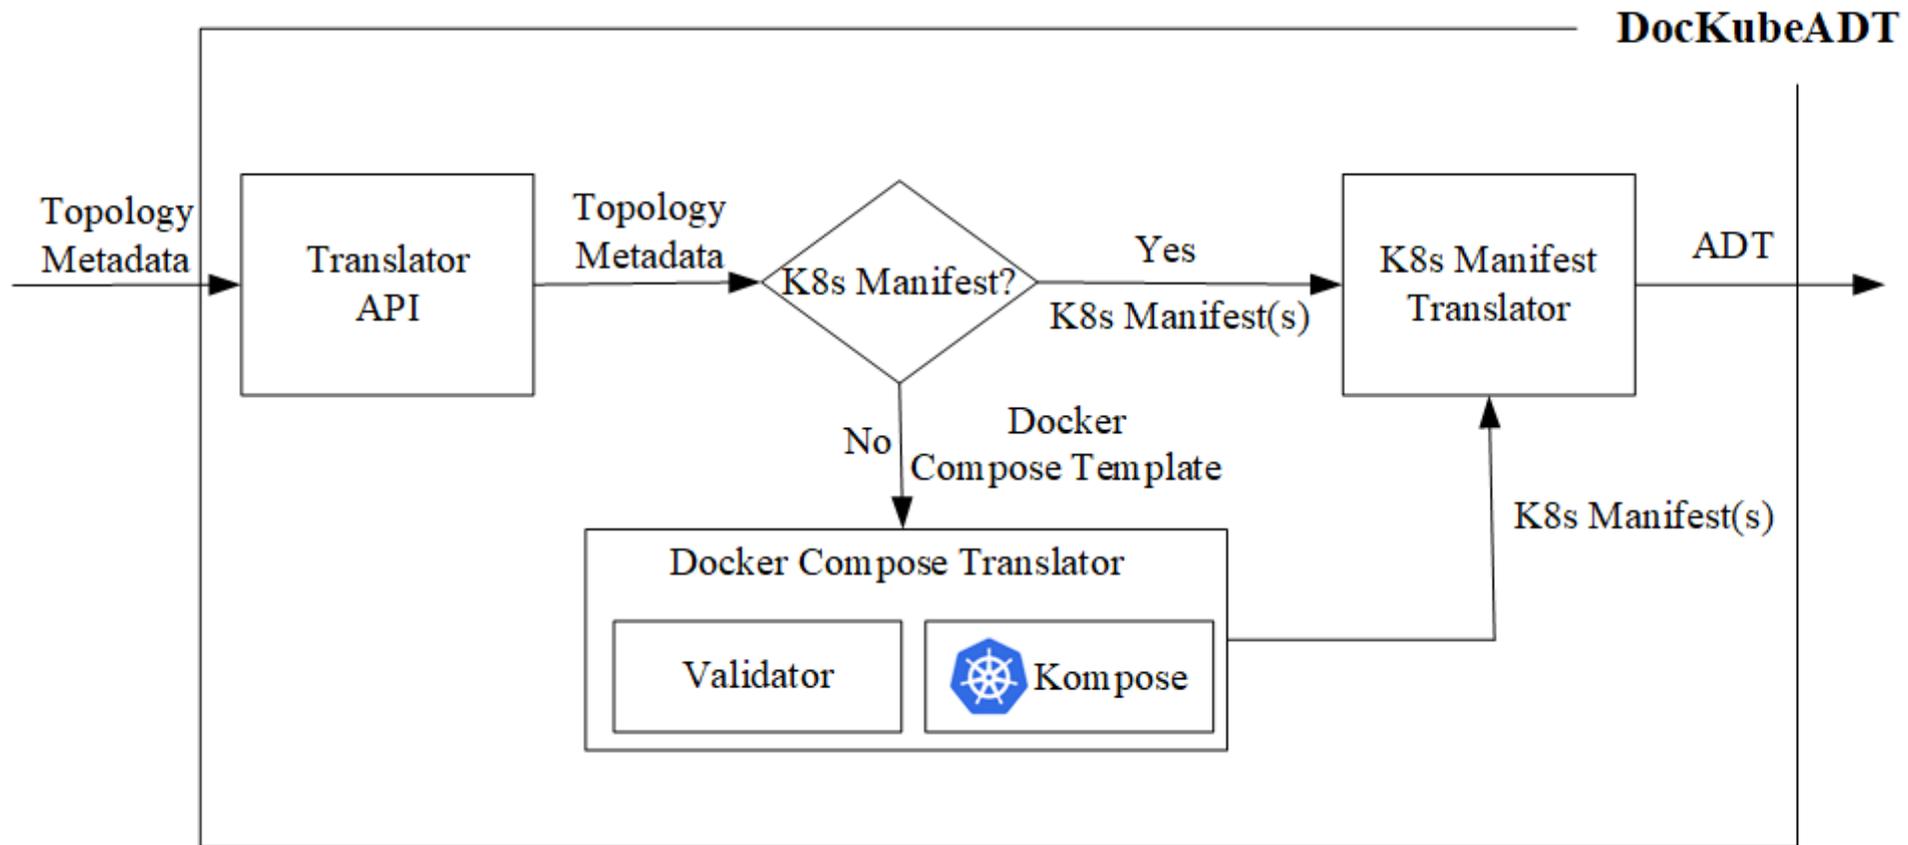
- ▮ all created by different authors at different phases



ADT generation in multiple steps



DocubeADT



How does it all work?



The screenshot shows the DIGITbrain website interface. At the top right, there are social media icons for @, LinkedIn, Twitter, and Facebook. Below these are navigation links: MARKETPLACE, ENTERPRISE, COMMUNITY, SUPPORT, and ACCOUNT (t.kiss@westmin...). The main heading is "DIGITbrain Solutions for Digital Twins". Below this is a paragraph: "DIGITbrain is the management platform for advance manufacturing technologies. It enables your company to boost innovation on demand through the power of advanced simulation, modeling, and data analytics for the manufacturing industry." At the bottom of this section is a "SIGN UP" button.

The diagram, titled "PURE DIGITAL LIFESPAN" at the top and "PHYSICAL & DIGITAL TWIN LIFESPAN" at the bottom, illustrates a circular flow of data and processes. The central element is a brain icon. Surrounding it are boxes for "DATA", "MODELS", "RESOURCES", and "ALGORITHMS". The outer ring consists of boxes for "ENGINEERING", "PRODUCTION", "OPERATION", "MRO", "RECYCLING", and "DISTRIBUTION". Arrows indicate a clockwise flow between these stages, with bidirectional arrows connecting the central data/model boxes to the outer process boxes.



Publish assets, build and execute DMA tuples



The screenshot shows the Digitbrain Assets marketplace. At the top left is the Digitbrain logo. To its right are navigation links: MARKETPLACE, ENTERPRISE, COMMUNITY, SUPPORT, and ACCOUNT (t.kiss@westmin...). Below this is a large teal banner with the word "Assets" in white, followed by the tagline "... optimizing your manufacturing processes with modular digital twins composed of expert assets...". Underneath the banner are filter options: MICROSERVICE (MS) [PUBLISHED], ALGORITHM (A), MODEL (M), and BEHAVIOR (M+A). On the right side of the filter area is a page indicator "1 - 16". Below the filters, three asset cards are visible: "RCclone 2" with a blue and dark blue logo, "RISTRA CPU 2" with the same blue and dark blue logo, and "CAELIA" with a black puzzle piece icon.



Publish assets, build and execute DMA tuples

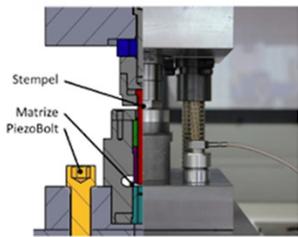


The screenshot shows the 'Create a new microservice' form in the DIGITBRAIN interface. The form is set against a light green background. At the top left is the DIGITBRAIN logo. To the right of the logo are navigation links: MARKETPLACE, ENTERPRISE, COMMUNITY, SUPPORT, and ACCOUNT (t.kiss@westmin...). The form itself is white and contains the following sections:

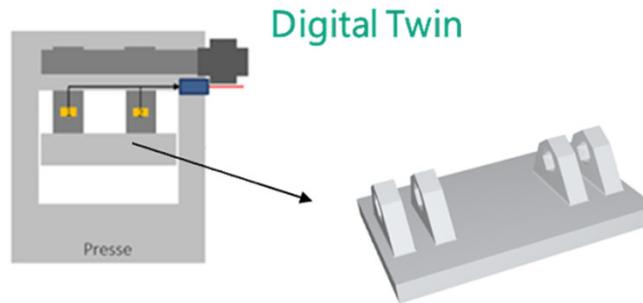
- Create a new microservice**
- Name**: A text input field with the placeholder 'Name'. Below it, a note states 'Maximum 64 characters.'
- Icon**: A file selection area with a 'Choose File' button and the text 'No file chosen'. Below it, a note states 'Recommended size: 200 x 200 pixels.' and a small box says 'No file chosen'.
- Short description**: A text input field with the placeholder 'Short description'. Below it, a note states 'Maximum 256 characters.'
- Techniques**: A list box containing 'Additive Manufacturing' and 'Artificial Intelligence'.



Example: Digital Twin of a Punching Machine



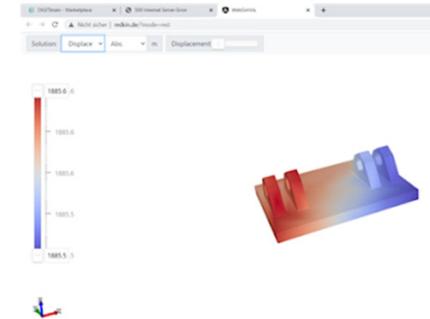
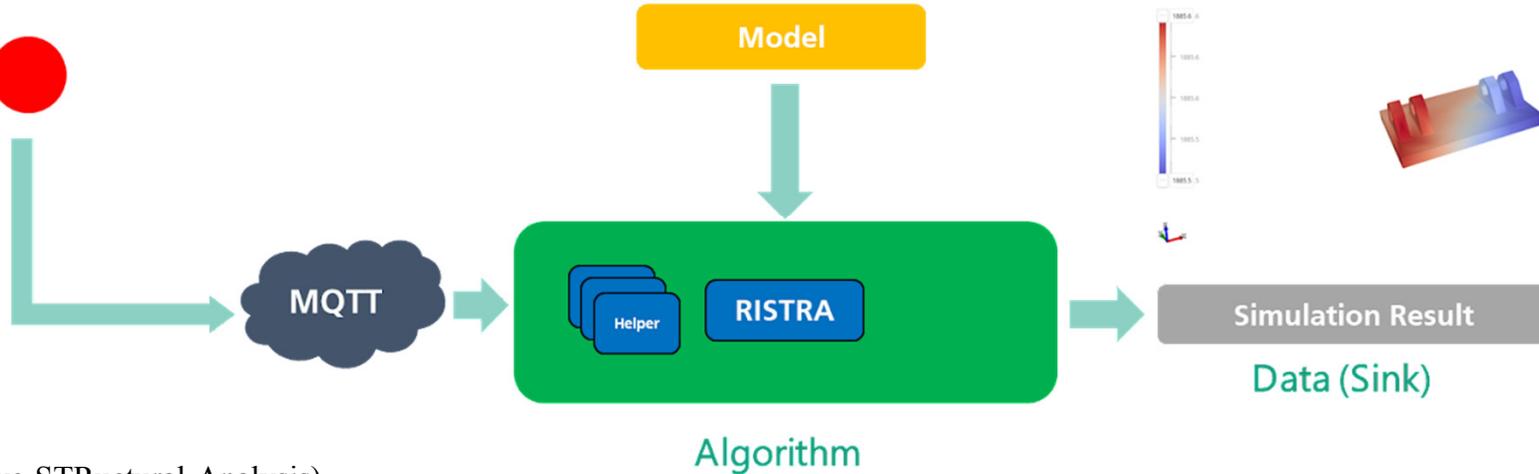
Industrial Product



Digital Twin

- Data
- Microservice
- Algorithm
- Model

Data (Source) ●



Simulation Result
Data (Sink)

RISTRA (Rapid Interactive STRuctural Analysis), a software library for fast structural analysis.

Example provided Fraunhofer IGD – Andre Stork, Daniel Weber, Johannes Mueller-Roemer, Lars Lotter, Maxim Redkin



How to build the DMA Tuple?



Model: Created with RISTRA using a CAD model



MODEL.FILENAME: press.zip

MODEL.PATH: /Fraunhofer/data

MODEL.URI: https://plg-cyfronet-01.datahub.egi.eu



MODEL_ID_Press

Microservices: Core and additional functionality

RISTRA
Operates on the Model

MS_ID_RISTRA

Model Retriever

MS_ID_MODELRETRIEVER

Data Bridge

MS_ID_DATABRIDGE

Algorithm: Includes three Microservices

Algorithm: RISTRA

MS_ID_RISTRA

MS_ID_MODELRETRIEVER

MS_ID_DATABRIDGE



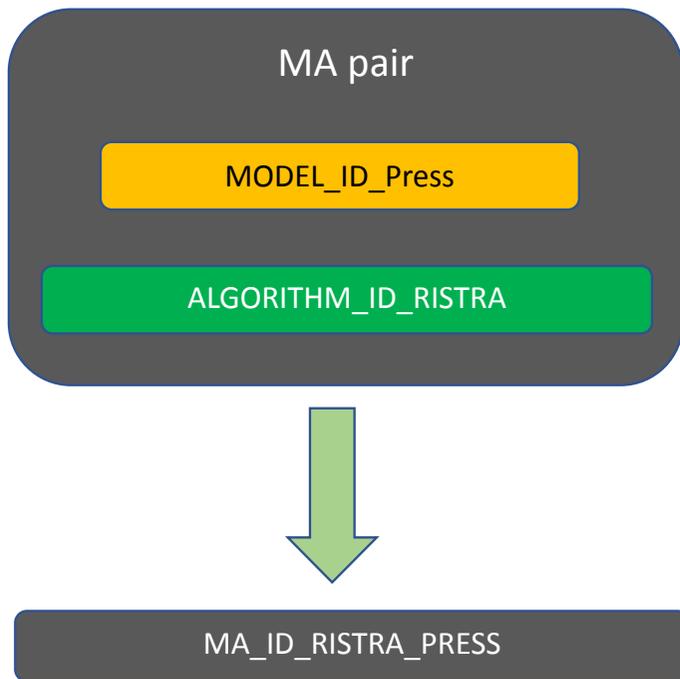
ALGORITHM_ID_RISTRA



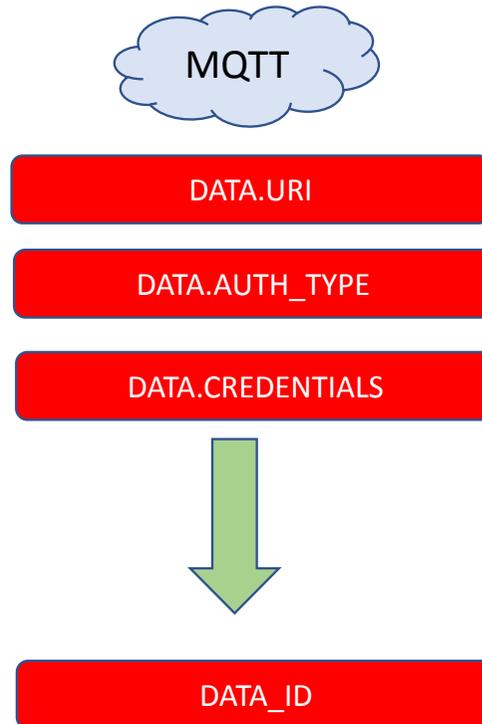
How to build the DMA Tuple?



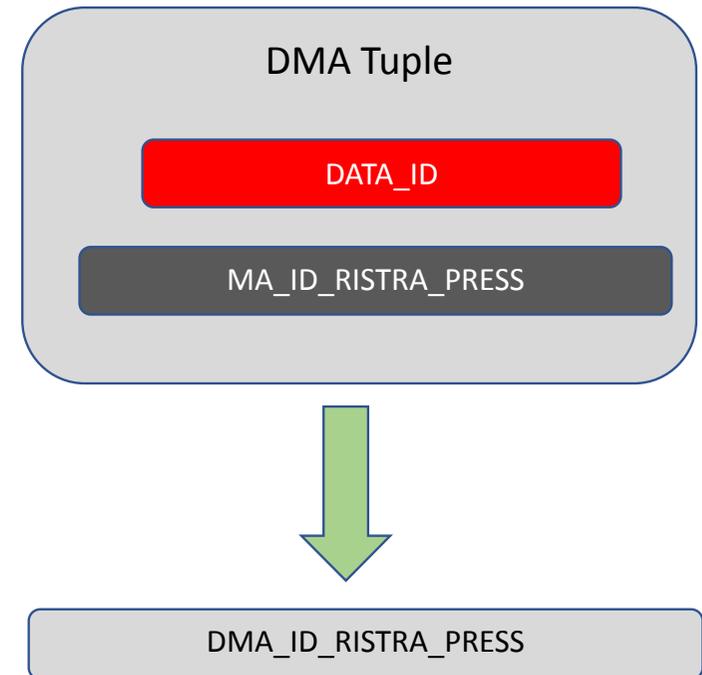
MA Pair: Describes the behaviour of an IP Family



Data source: From IoT sensors via MQTT



DMA Tuple: The executable Digital Twin



How to execute the DMA Tuple?



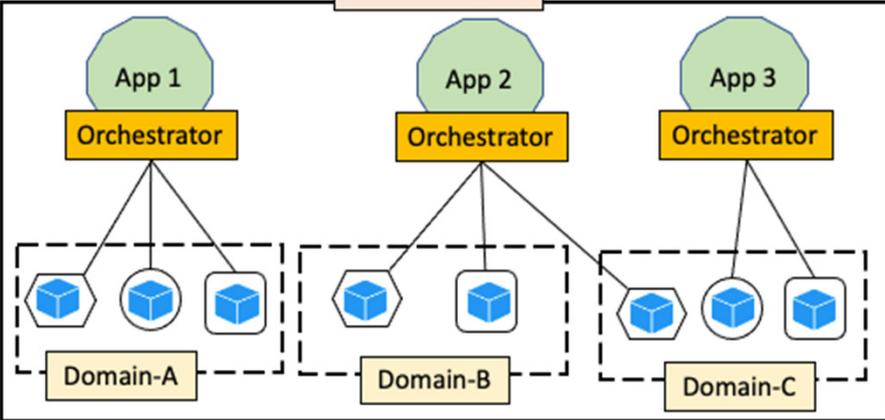
- Once DMA Tuple is built ADT generator can create TOSCA ADT for execution
- Copy of MiCADO is deployed by the platform
- ADT is passed to MiCADO to execute
- Resources are allocated based on ADT - RISTRA core simulation executes on edge



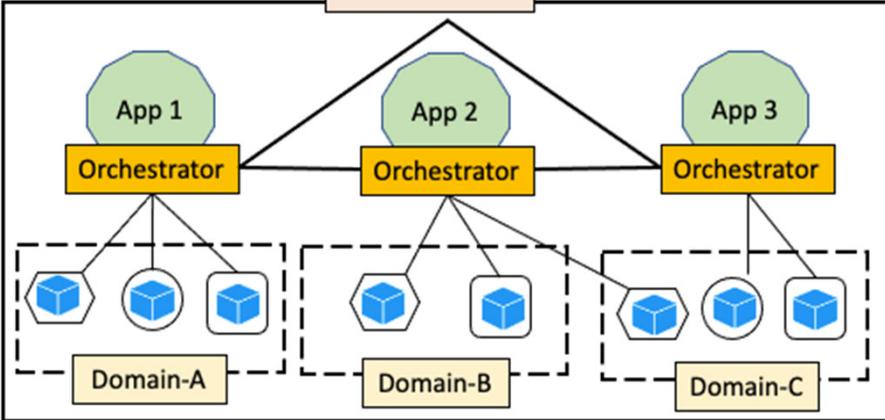
Where to go next? – Decentralised orchestration



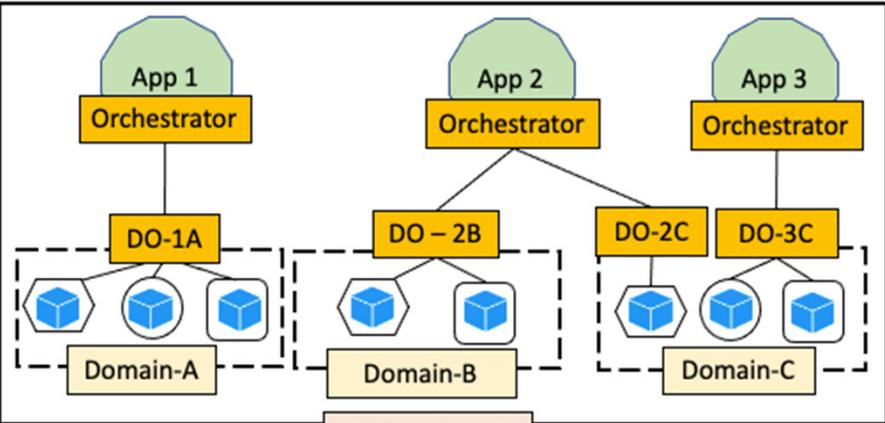
Scenario - A



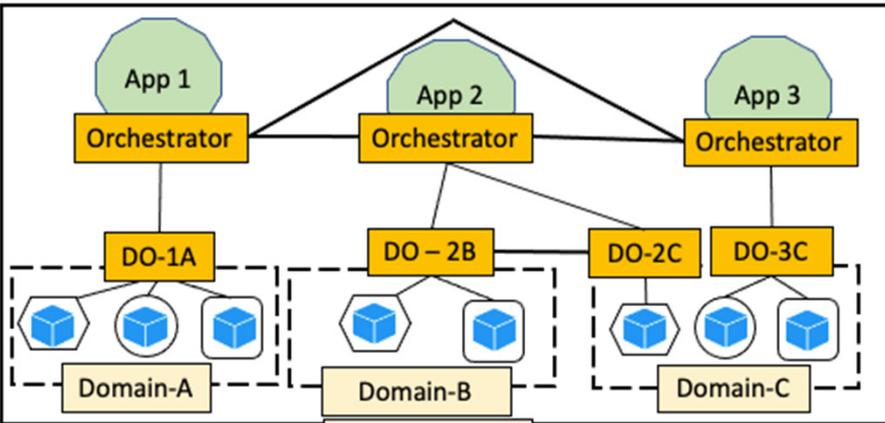
Scenario - B



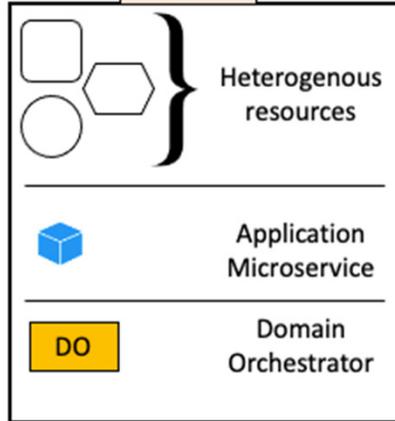
Scenario - C



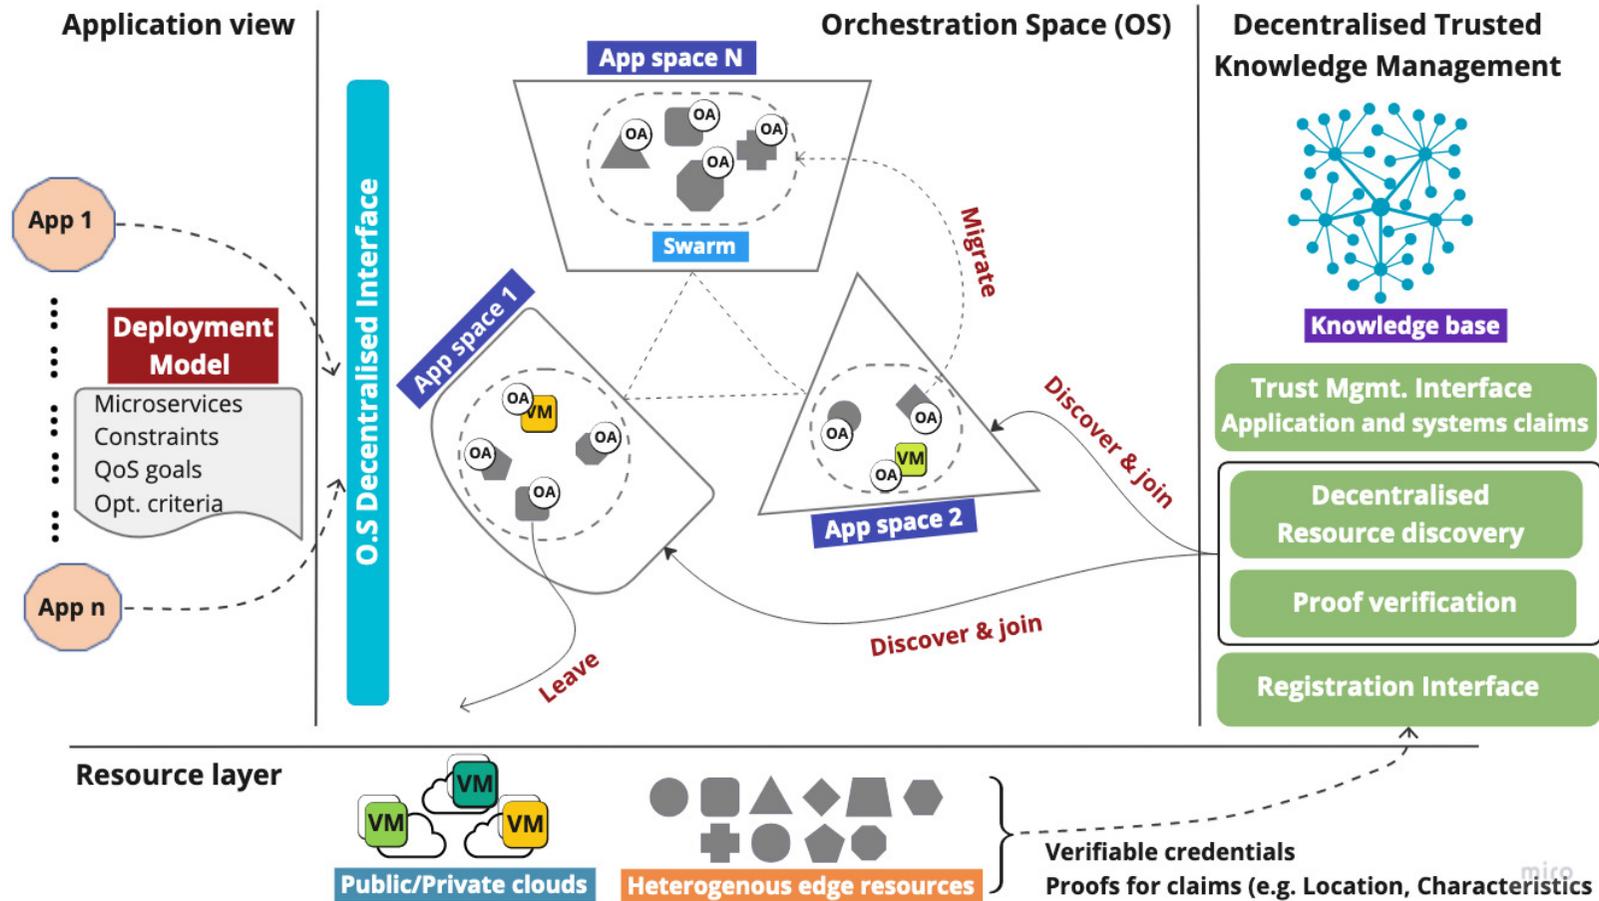
Scenario - D



Legend



Where to go next? - AI-Driven Swarm-based orchestration



The current CPC team behind all of this



Tamas Kiss



Gabor Terstyanszky



Gabriele Pierantoni



Huseyin Dagdeviren



Francesco Tusa



Hamed Hamzeh



James DesLauriers



Antonis Michalas



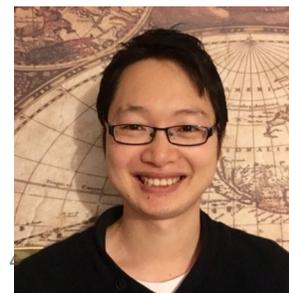
Amjad Ullah



Jozsef Kovacs



Alim Gias



Huankai Chen



Yang Ma



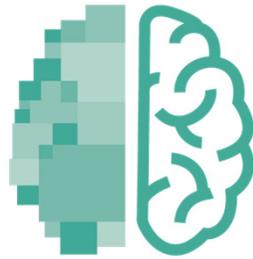
Dimitris Kagialis



David Chan You Fee

DIGITbrain has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 952071





DIGITBRAIN

Thank you for your attention!

Prof Dr Tamas Kiss

University of Westminster



t.kiss@westminster.ac.uk

 www.digitbrain.eu

 www.linkedin.com/groups/12439191

 www.twitter.com/digitbrain_eu

 www.facebook.com/DIGITbrainProject

DIGITbrain has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 952071

