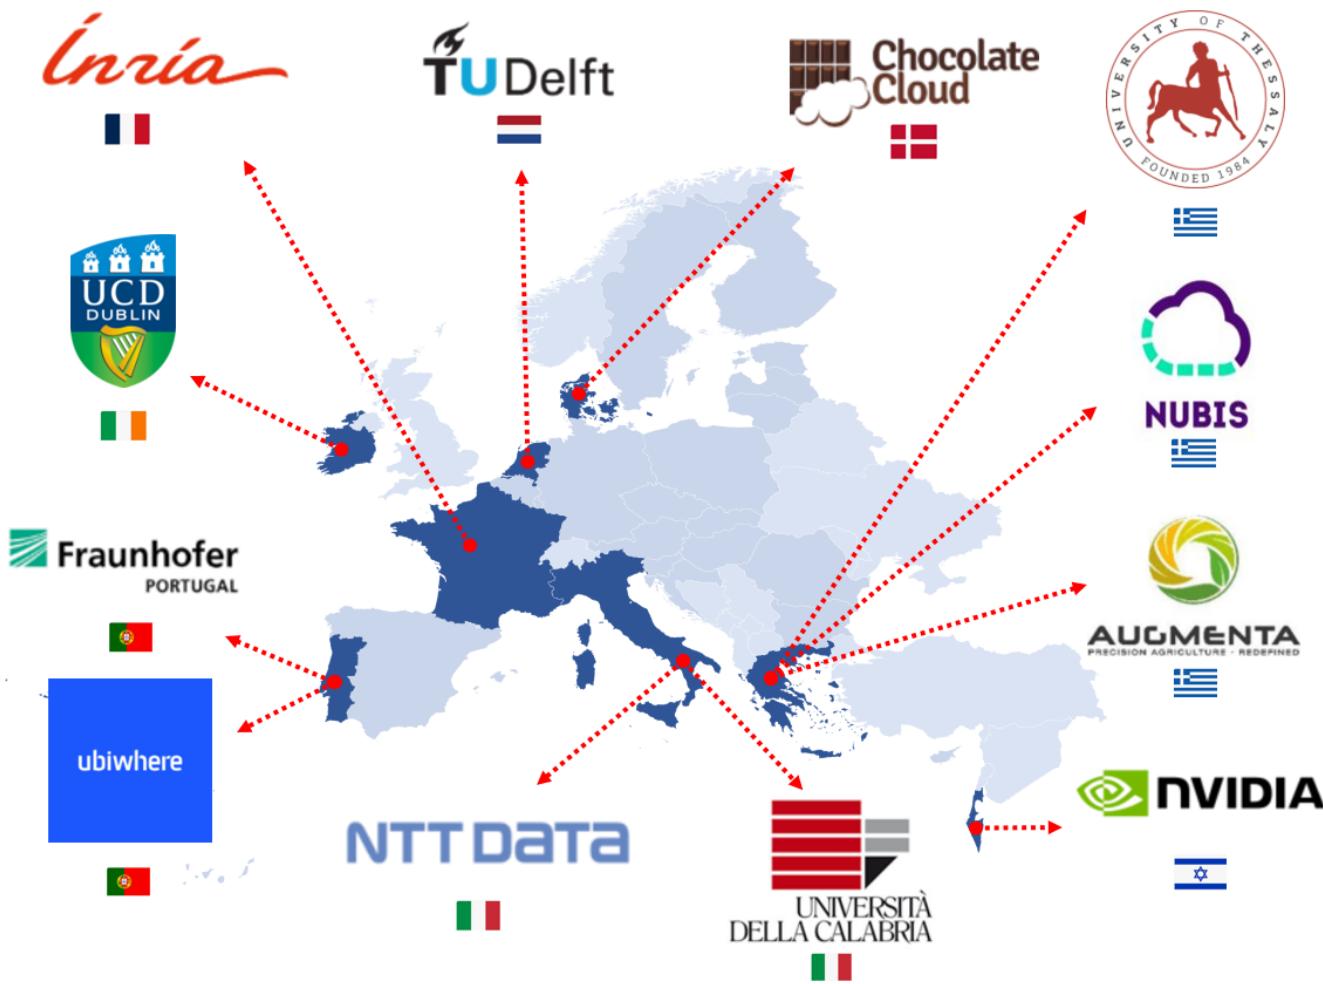


# MLSysOps

# Machine Learning for Autonomic System Operation in the Heterogeneous Edge-Cloud Continuum

# Nikolaos Bellas

## University of Thessaly



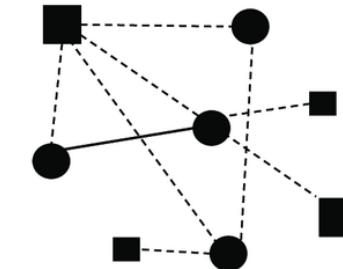
# Trends

Applications have moved outside the cloud

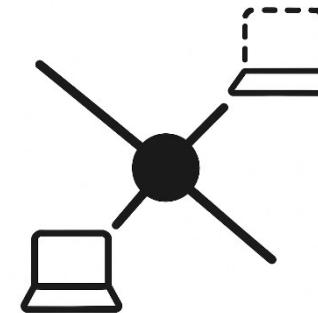
- Involving nodes and resources toward or directly at the edge of the Internet...
- ... and including powerful but also resource-constrained IoT devices



Scale



Heterogeneity



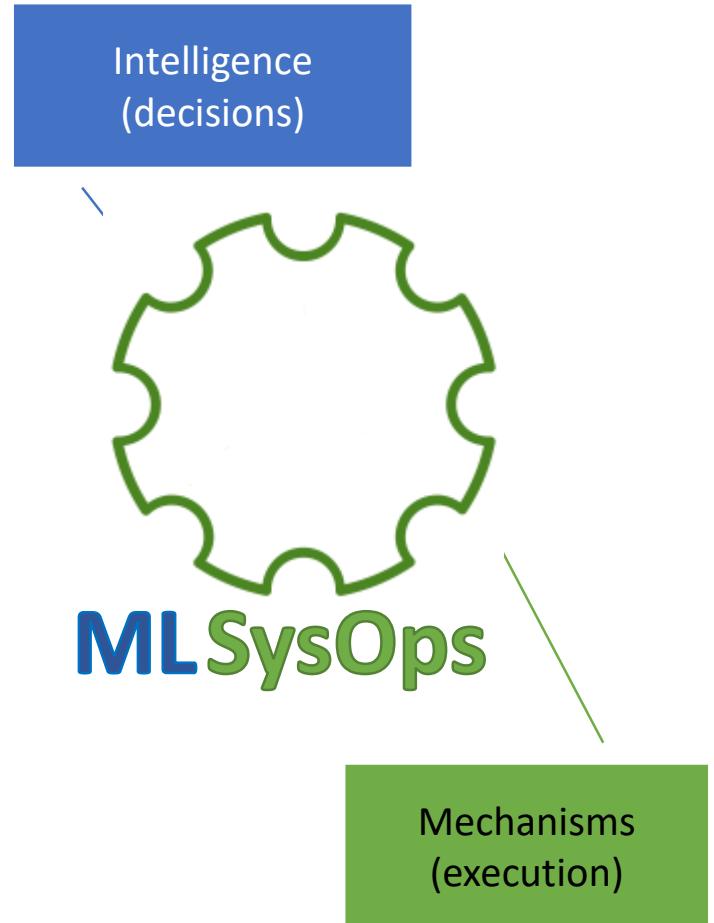
Volatility

**Impossible to monitor/manage large systems by rule-based approaches**

# The MLSysOps project

*Autonomic system management and configuration  
in the **cloud-edge-IoT continuum** using AI/ML methods*

- Focus on modular, distributed applications
  - Comprised out of independently executable components (containers)
- Explore different management aspects
  - Deployment, computing, storage, communication/networking, trust
- Disassociate management from control
  - Develop AI/ML-ready (policy-neutral) mechanisms
  - Take decisions using suitable ML models
- Key AI/ML properties
  - Distributed / hierarchical approach
  - Continual learning / efficient model retraining
  - Explainability
- Extensive data generation evaluation via real-world application testbeds, research testbeds and simulators



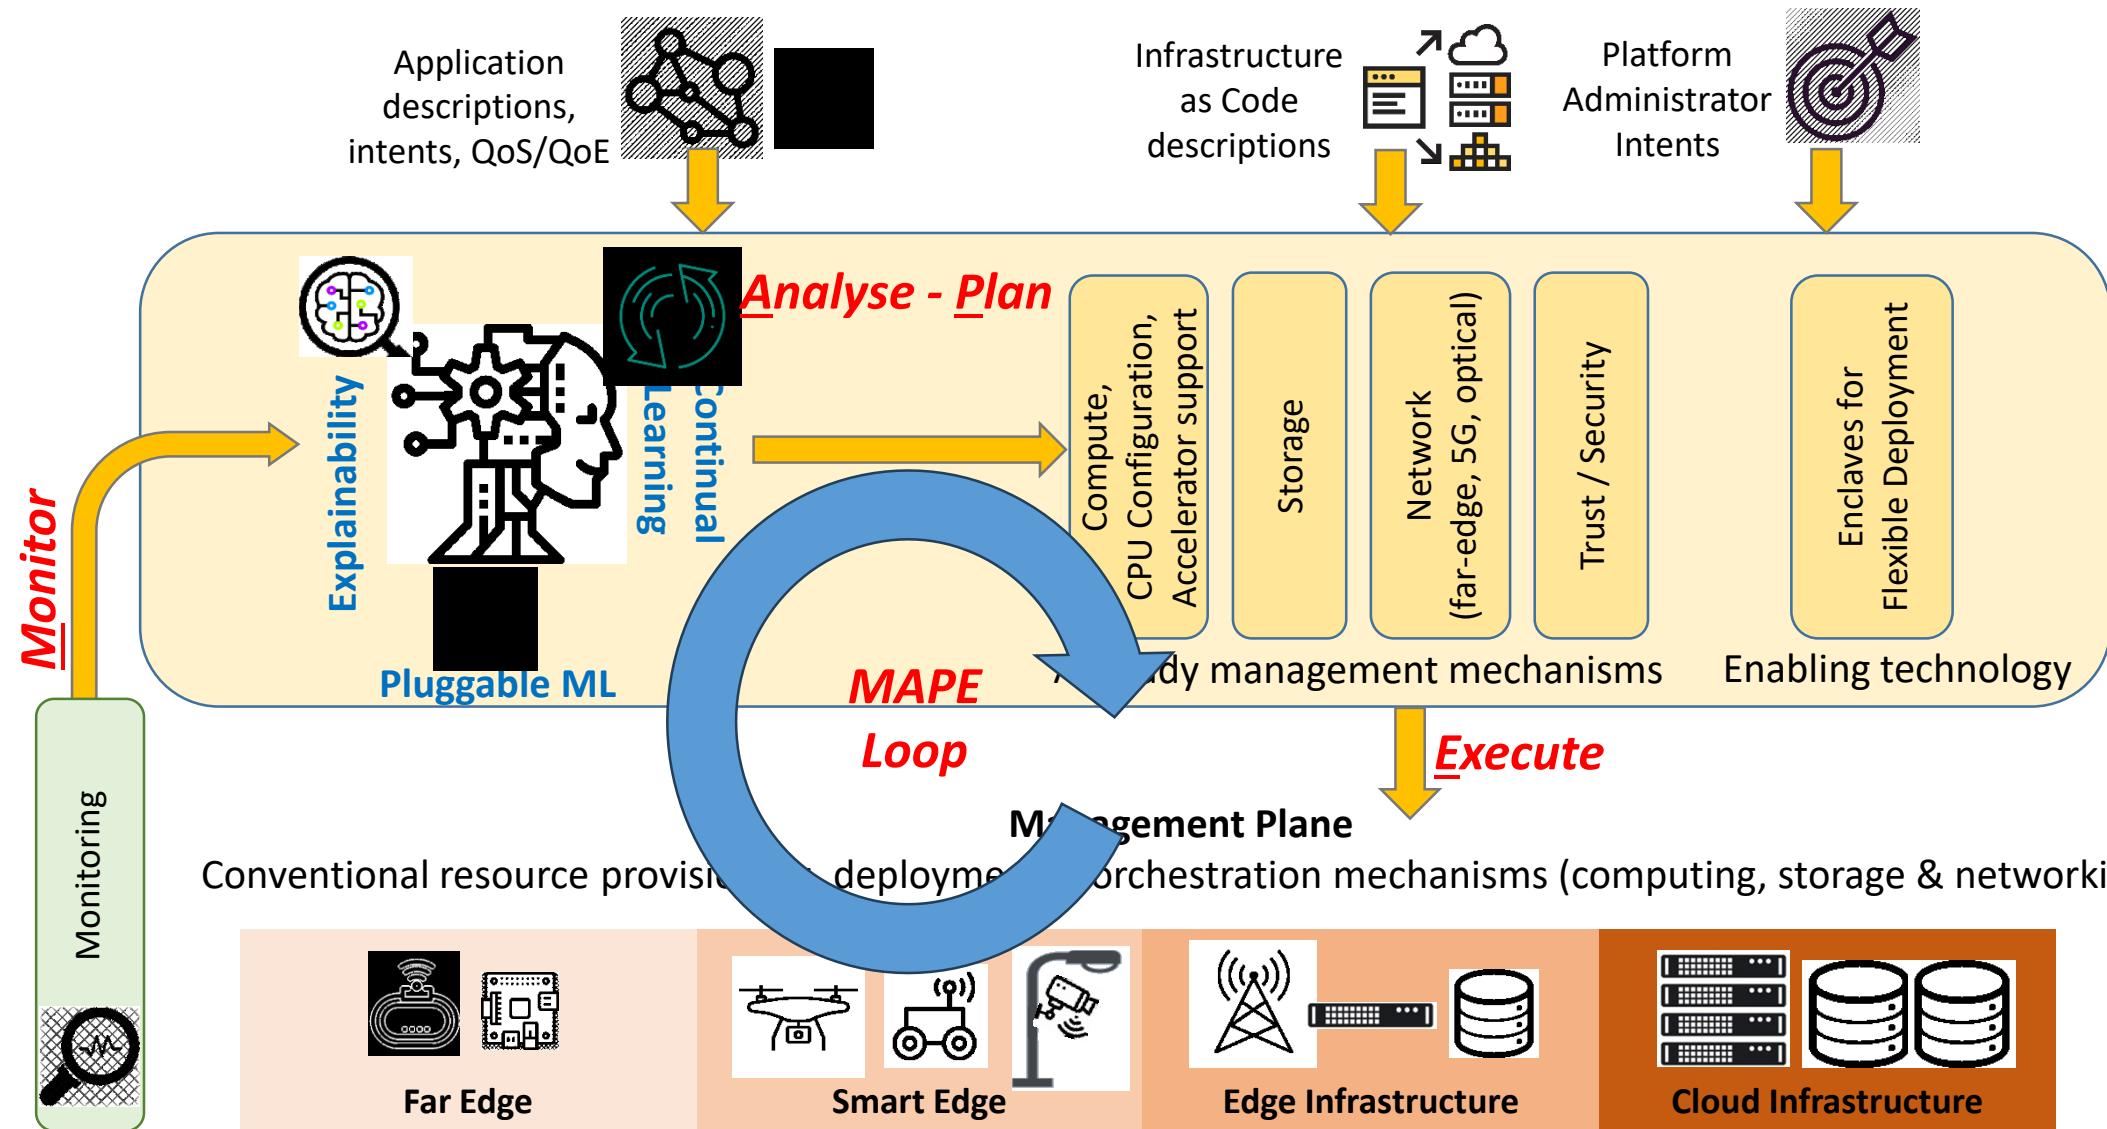


European  
Commission

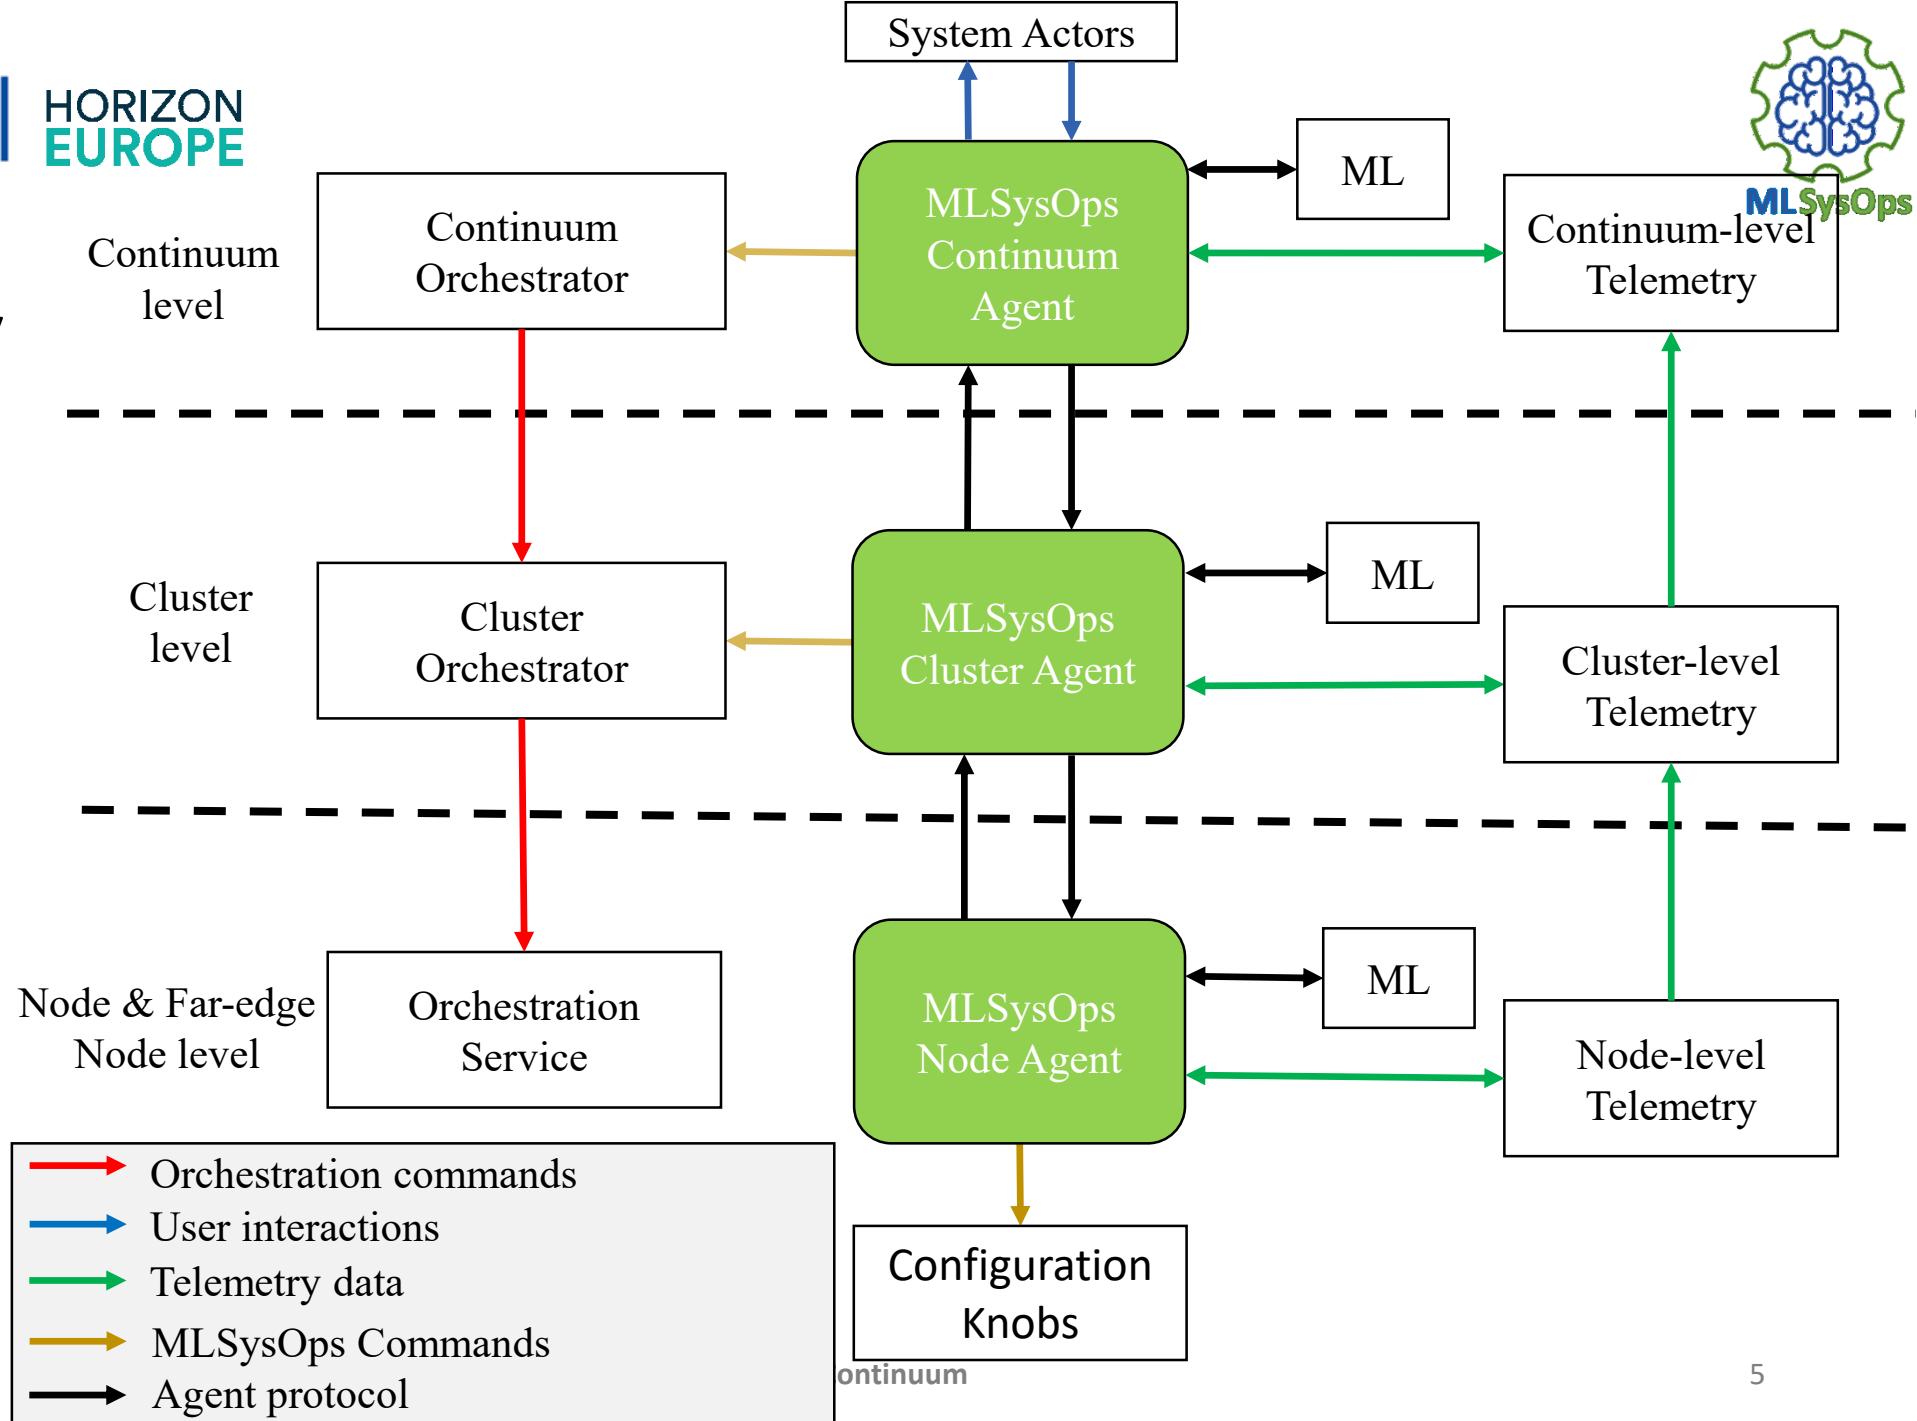
HORIZON  
EUROPE



# How? Concept



# How? Hierarchy



# Two real-world application use cases

## Smart Cities



Improve existing system via AI/ML-driven control to manage/configure application modules that can be deployed on smart lampposts and/or to datacenters.

## Smart Agriculture



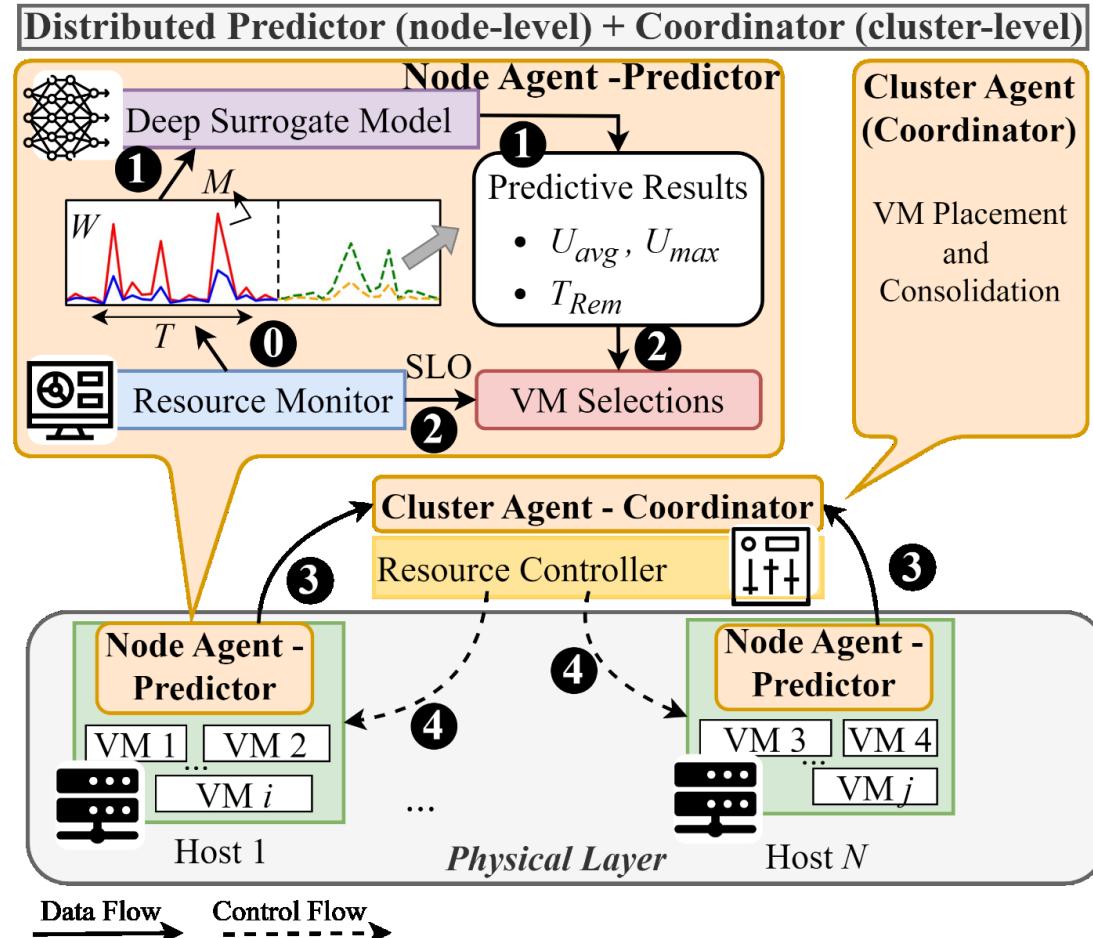
Improve existing system via AI/ML-driven control to jointly manage/configure the image processing application pipeline on a tractor and a companion drone.

# One of the many outcomes

Can we do better with ML-based resource management for the Cloud?

# Solution Overview

- Node-level ML model
  - Runs locally on each node
  - Predicts future CPU utilization and VM lifetime
  - Drives VM management decisions
- Cluster agent
  - Receives predictor output
  - Applies a heuristic to globally optimize cluster resources



# Results

- Migration counts
  - **Decrease of 41.33%** compared with SoTA conventional (non-ML) policy
  - Only 7.5% over oracle (Ground Truth)
- SLO violations
  - **Lower by 34.98%** compared with SoTA conventional policy
  - A mere 1.1% more than oracle (Ground Truth)

# Coming soon...

- Workshop organization
  - 2<sup>nd</sup> ML4ECS @ HiPEAC 2026 (Jan. 26, 2026)
  - In collaboration with other projects of the cluster  
<https://ml4ecs.e-ce.uth.gr/>
- Hackathon organization
  - University of Calabria, Rende, Italy (Sep. 26-27, 2025)
  - Familiarize developers with the MLSysOps framework
  - Foster application & ML models development



# Open-Source Release

Open-source components of the MLSysOps framework

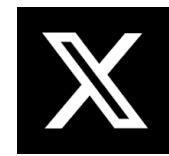
<https://github.com/mlsysops-eu/mlsysops-framework>

## Core Components

- Hierarchical Agent Architecture: Interfaces with orchestration/control systems and exposes an ML-model API for plug-and-play explainable/re-trainable models.
- Telemetry & Control Knobs: Collects metrics across the continuum and adjusts configuration (e.g., compute, network, storage, accelerator usage) dynamically.
- Distributed FaaS-style Executor: Enables function offloading across tiers to optimize latency, energy, and performance.
- Explainable ML & Reinforcement Learning Module: Offers transparent decisions, highlighting input factors influencing agent actions.
- Use-cases: Includes real applications focusing on smart cities and agriculture.



<https://mlsysops.eu>



@mlsysops



mlsysops



@mlsysopsproject

Thank  
you!

Newsletter

<https://mlsysops.eu/communication/>