European Commission | HORIZON EUROPE

MLSysOps

ML-Based Autonomic System Management in the Edge-Cloud Continuum

Christos Antonopoulos
Computer Systems Lab – ECE Department
University of Thessaly, Greece

MLSysOps Overview

HORIZON-CL4-2022-DATA-01-02
Cognitive Cloud: AI-enabled computing continuum from Cloud to Edge (RIA)

Grant ID: 101092912

12 partners

8 countries

1/1/2023-31/1/2026

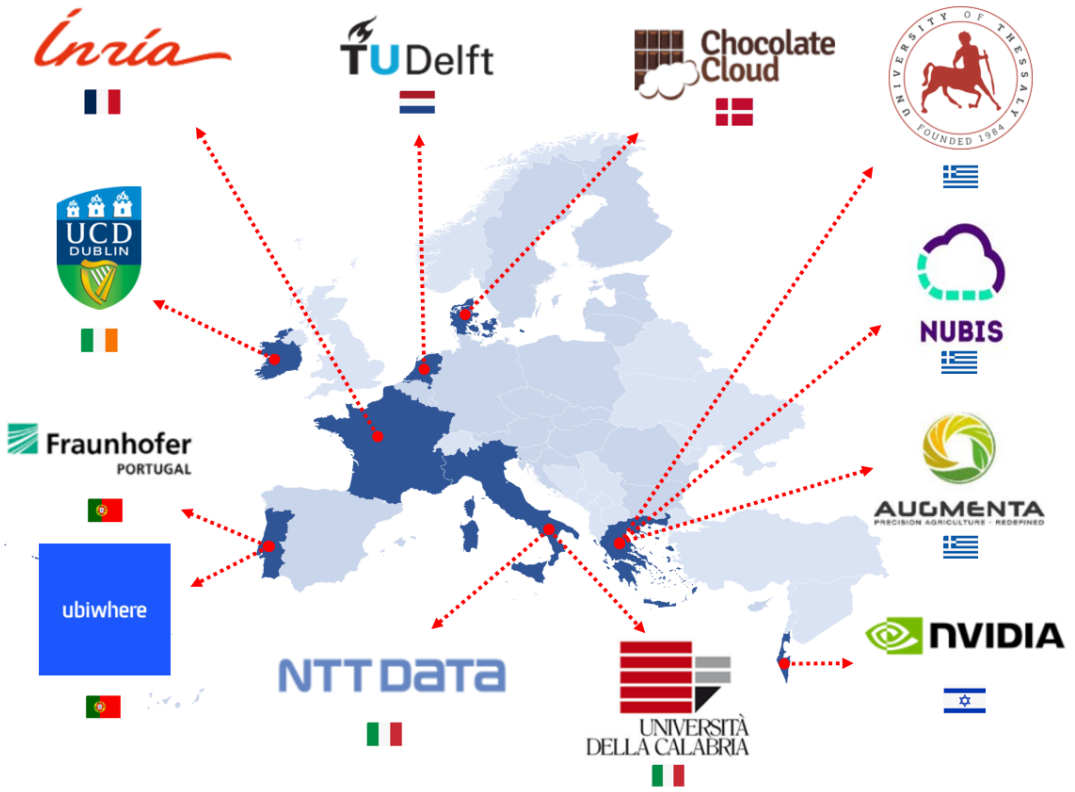

The trend…

Continuum systems

- Applications are moving outside the cloud
  - Start involving nodes and resources toward or directly at the edge of the Internet…
  - … including powerful but also resource-constrained IoT devices

The Challenges

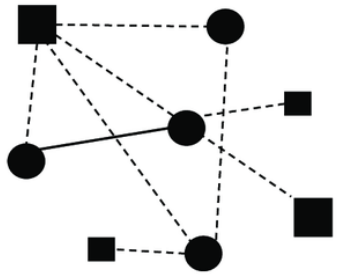

Heterogeneity

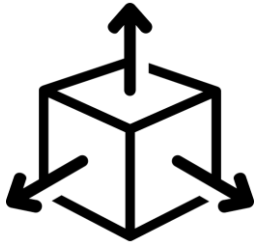

Scale

Security

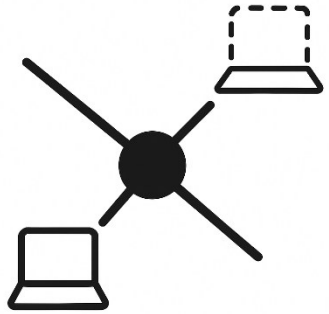

Volatility

Impossible to monitor/manage by a human

Are AI/ML methods for **autonomic system management** and configuration in the cloud-edge-IoT continuum **A FEASIBLE SOLUTION?**

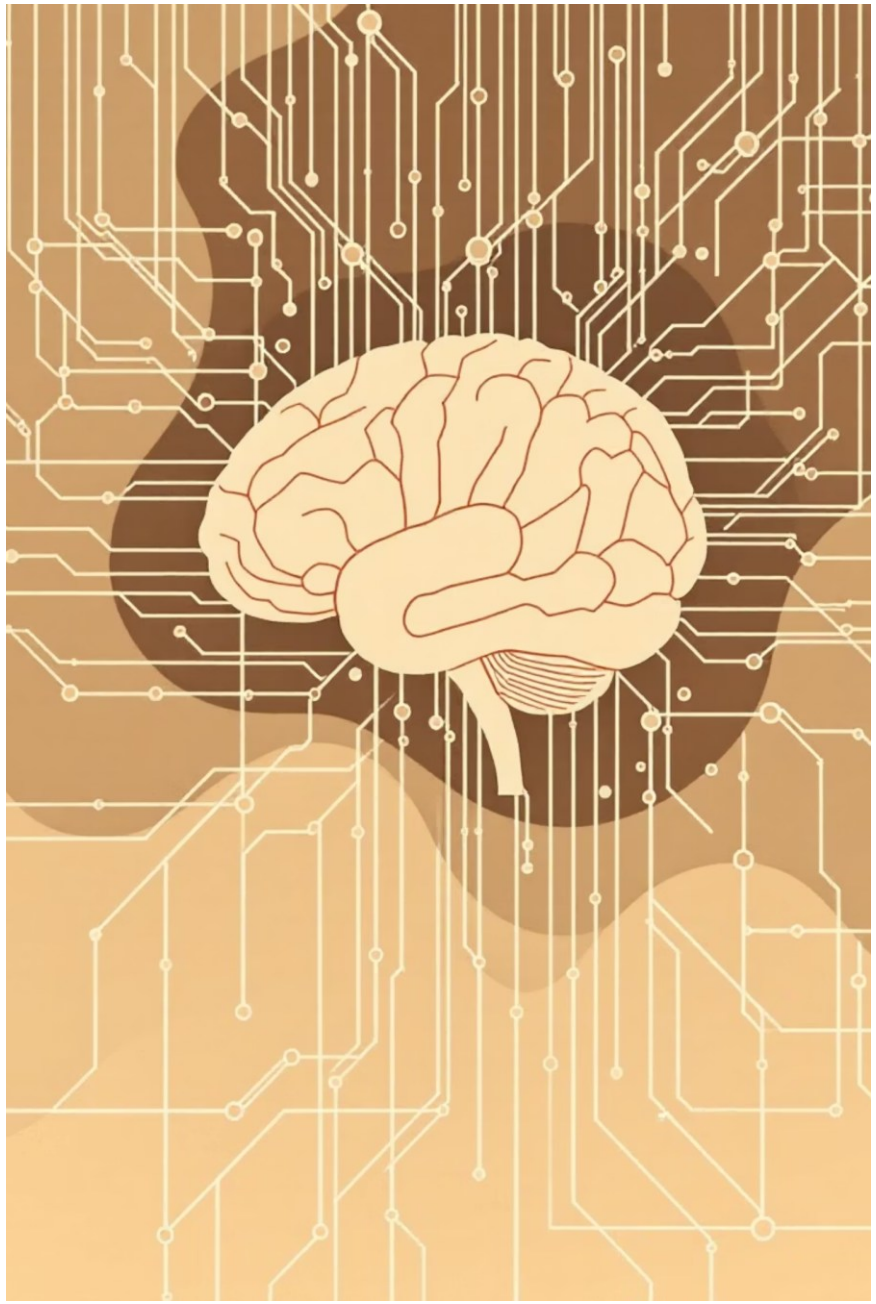

MLSysOps Concept

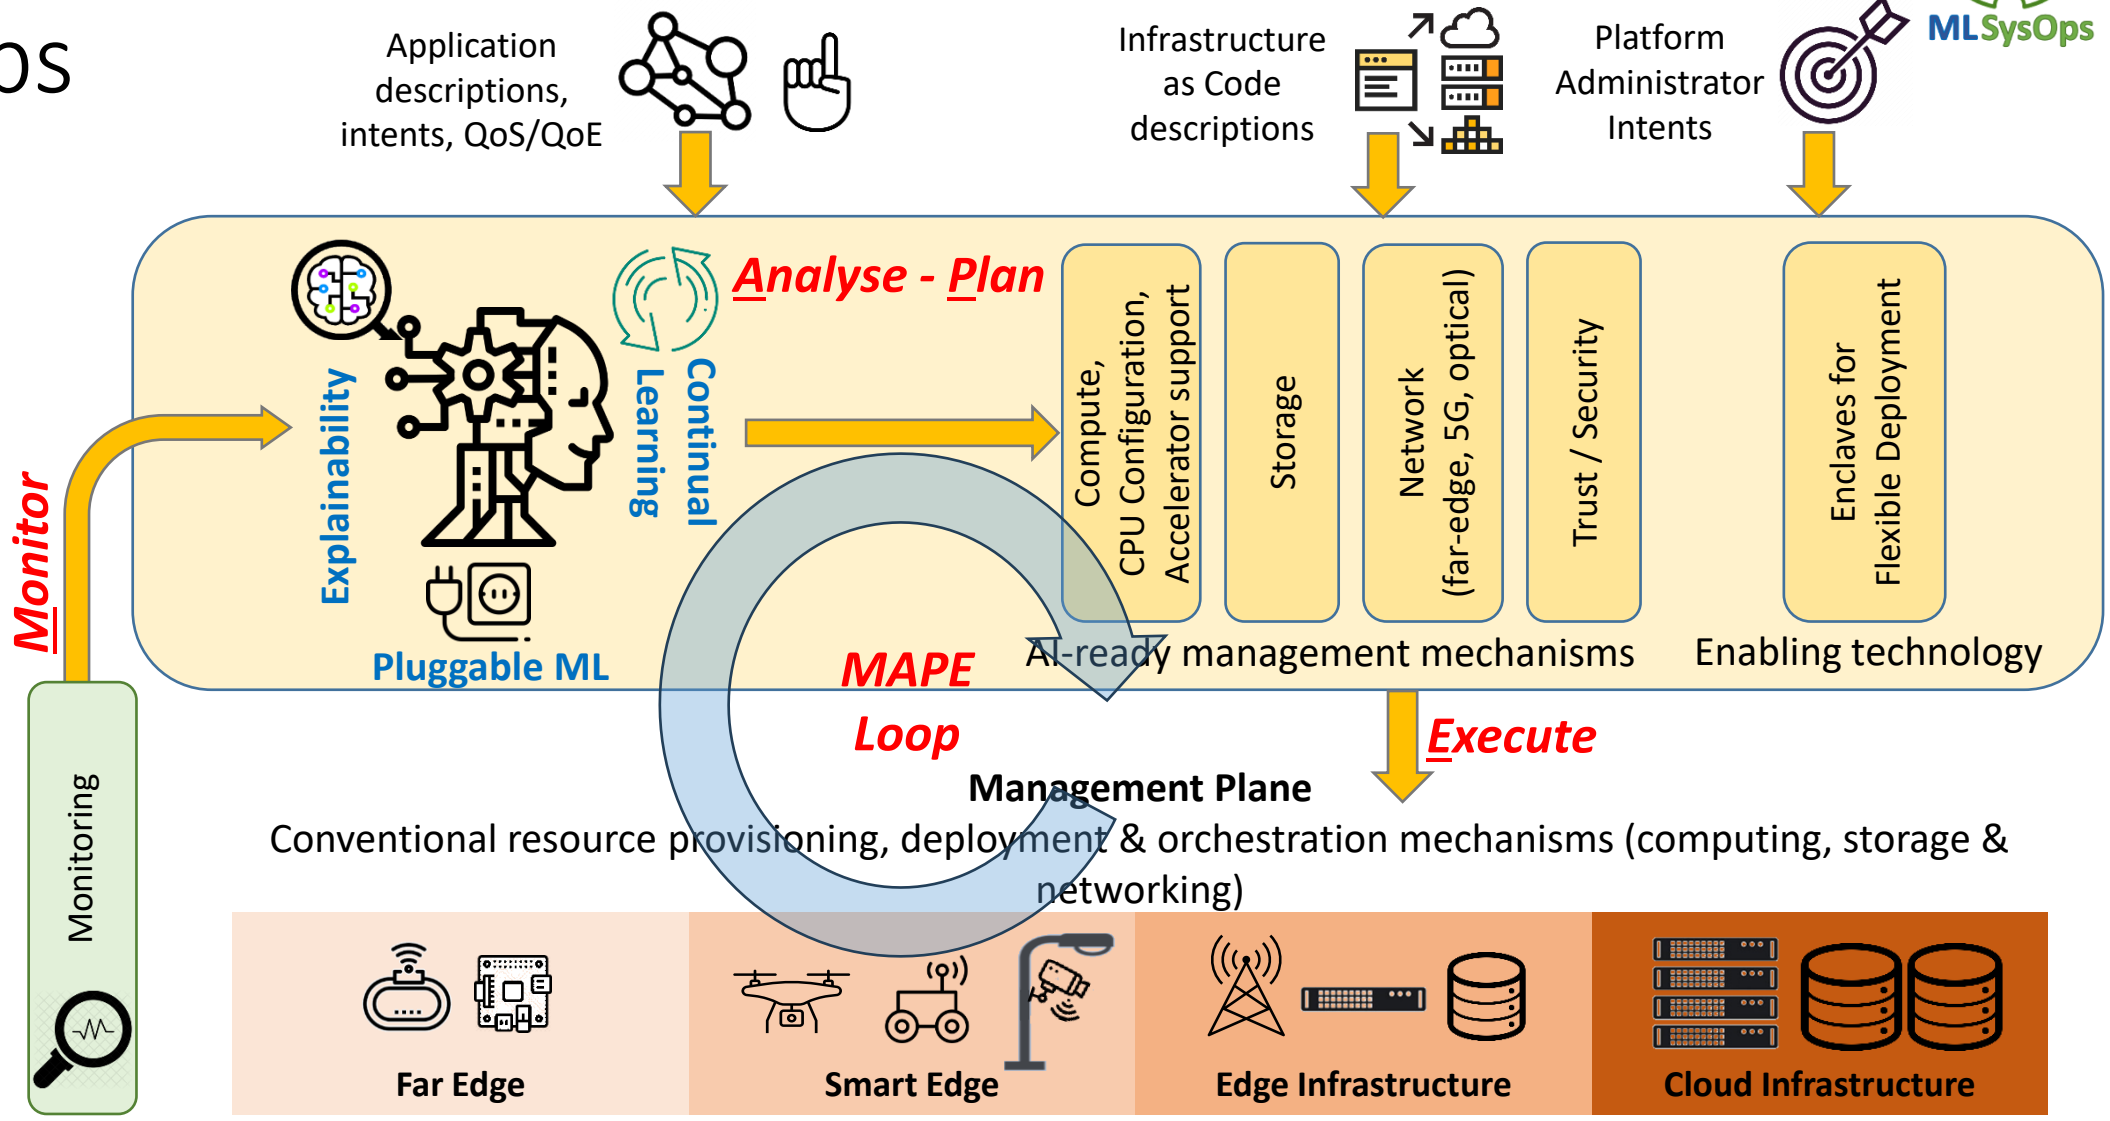

MLSysOps Approach

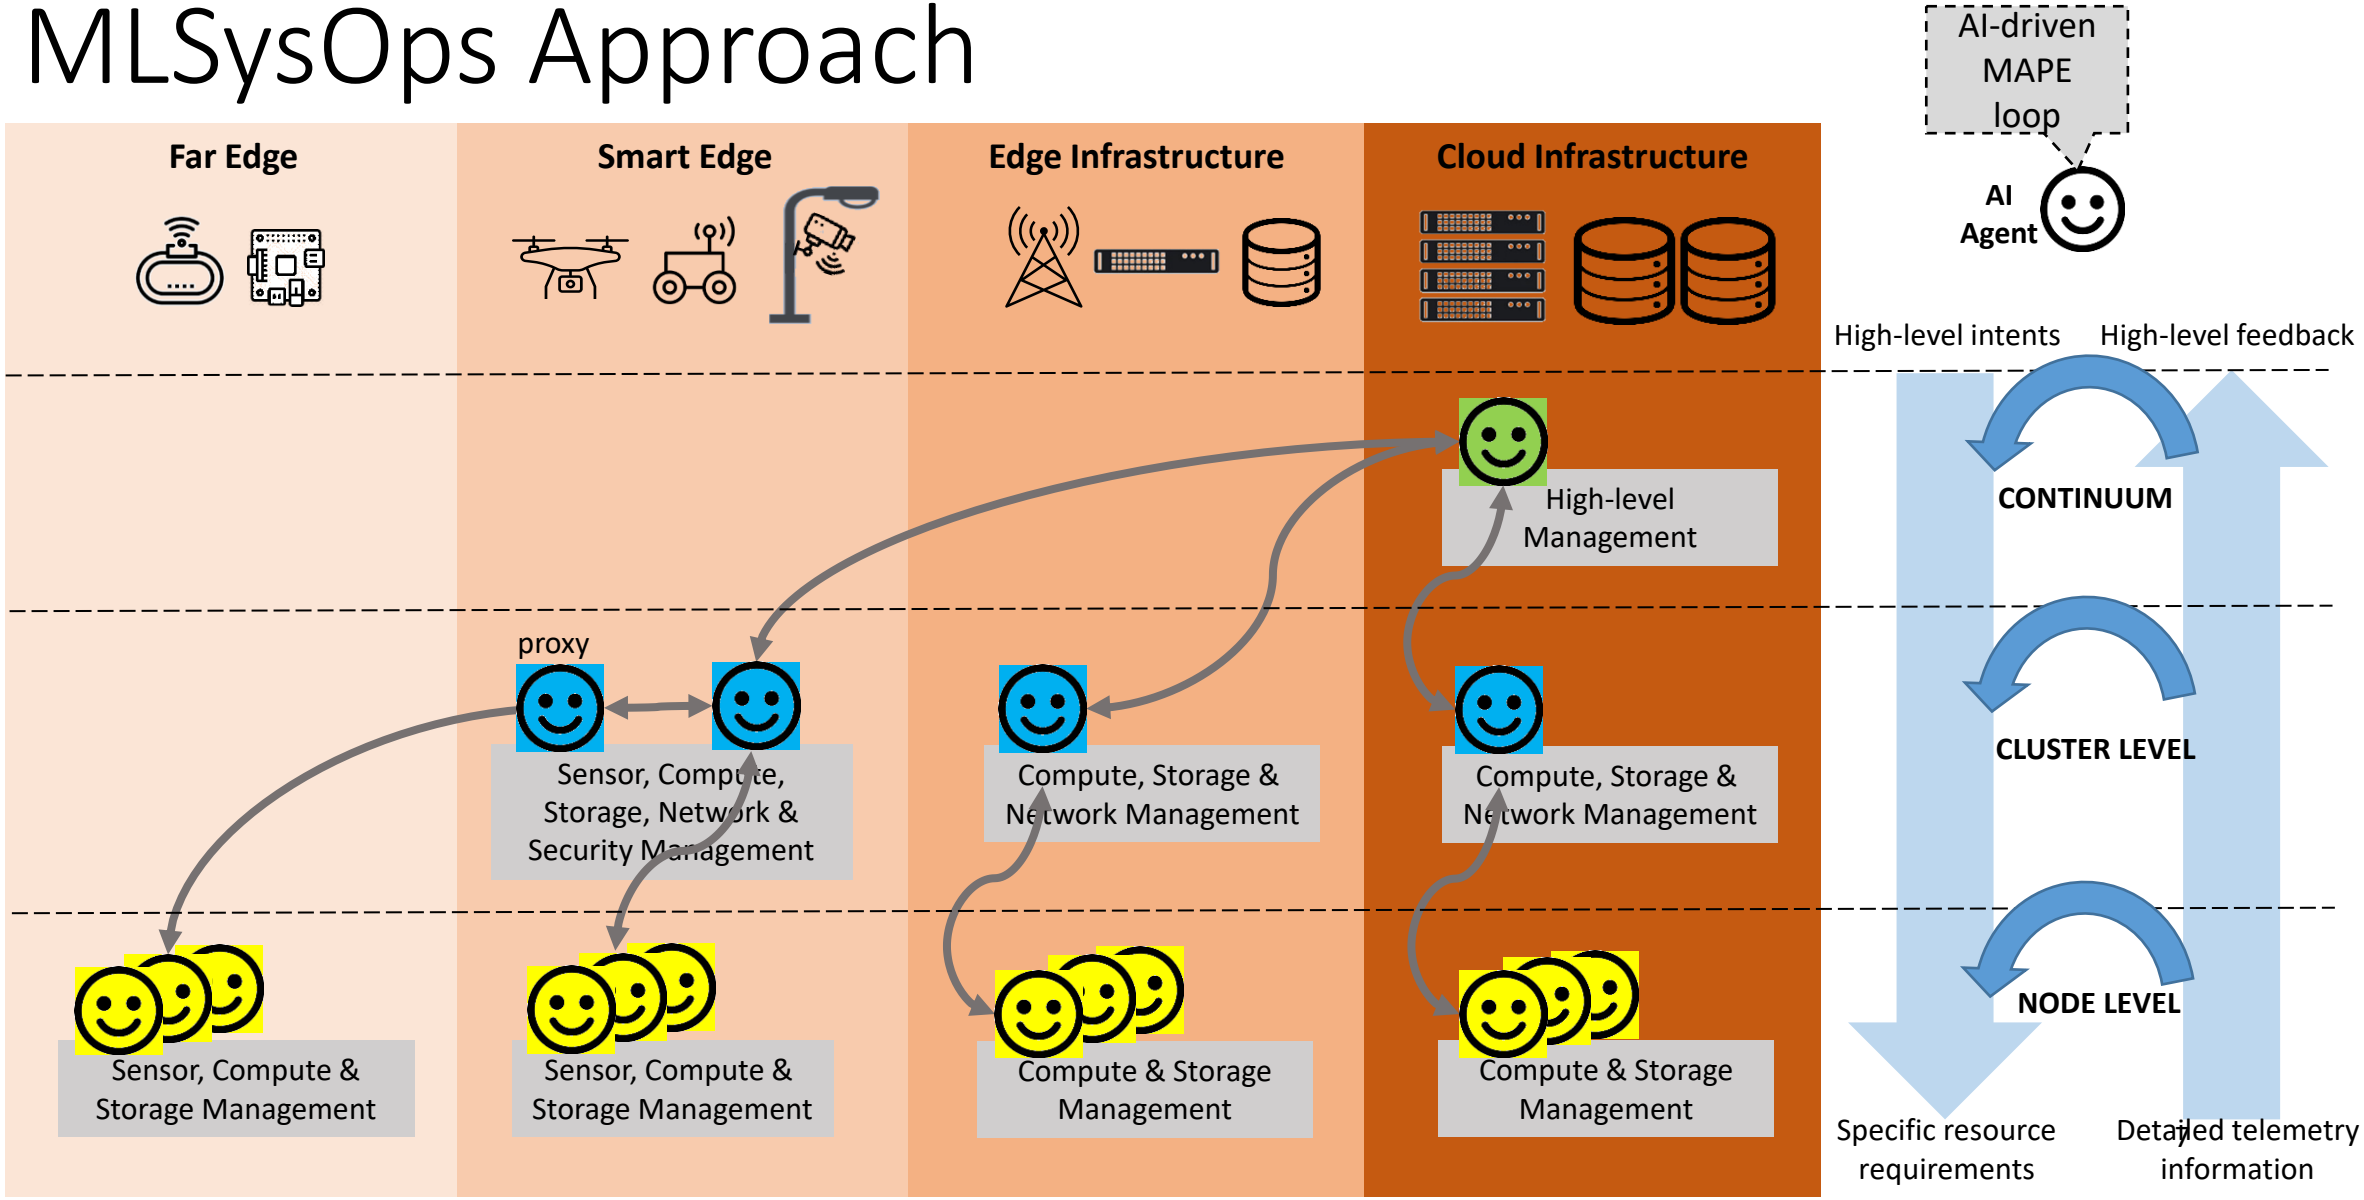

MLSysOps Instantiation

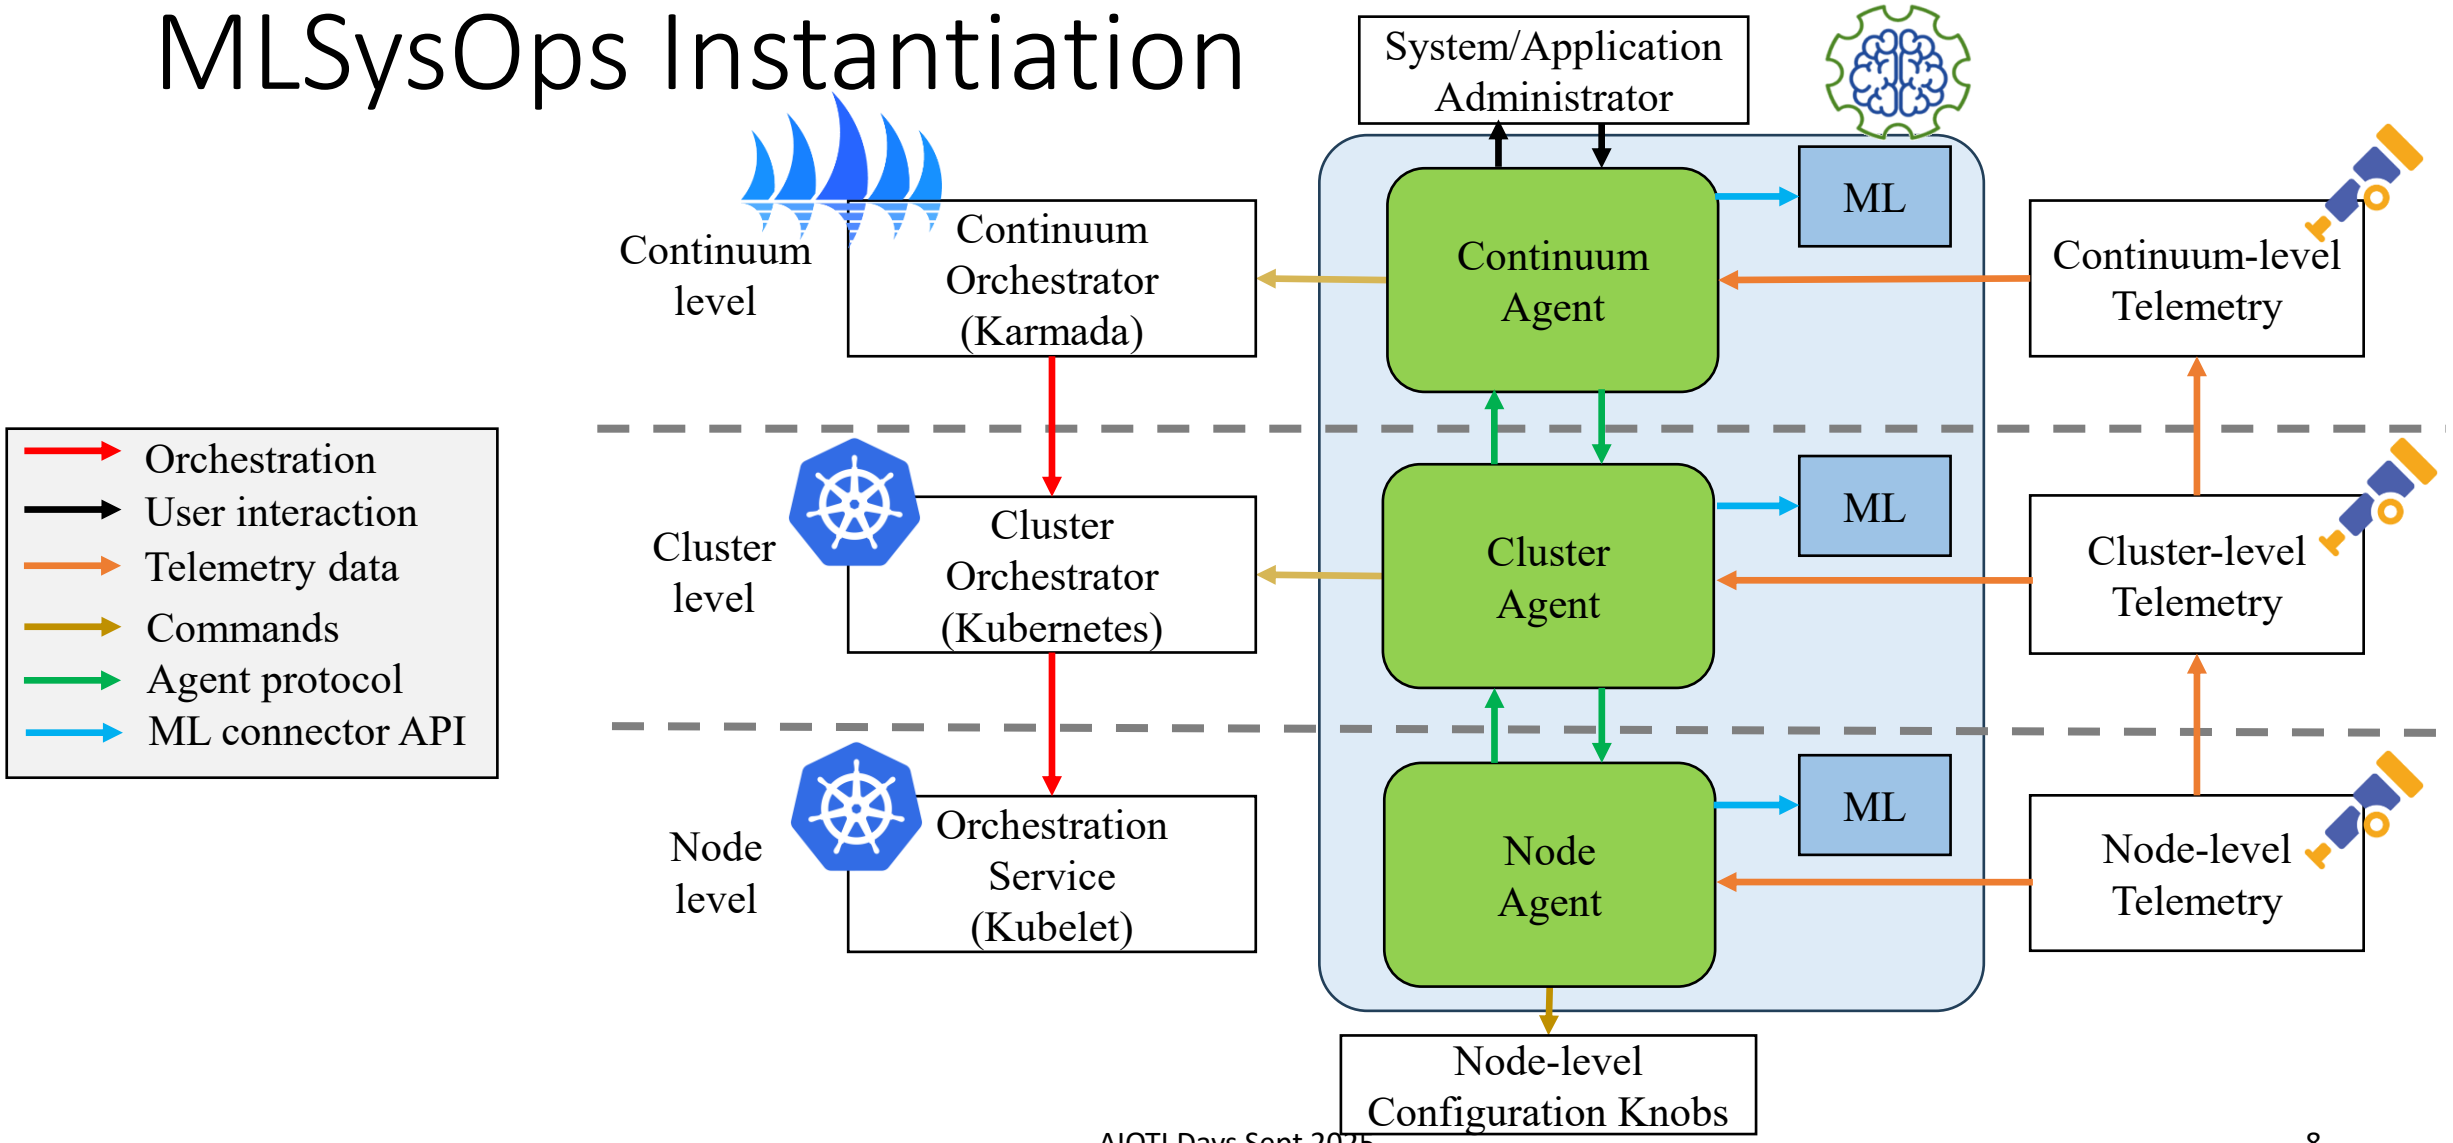

ML Models

→Reinforcement Learning

ML models may become outdated; adapt to a changing world/conditions

→Federated Learning (FL) & Split Federated Learning (SFL)

Privacy-preserving and distributed/scalable training

→Transfer Learning

Do not train from scratch for each new setup; reuse and adapt pre-trained models

→ML-as-Applications

ML models (and their training) can be deployed as special applications

These methodologies work synergistically to create adaptive, efficient, and privacy-preserving ML solutions for large-scale system management.

Use Cases

- **Smart Cities**

ubiwhere

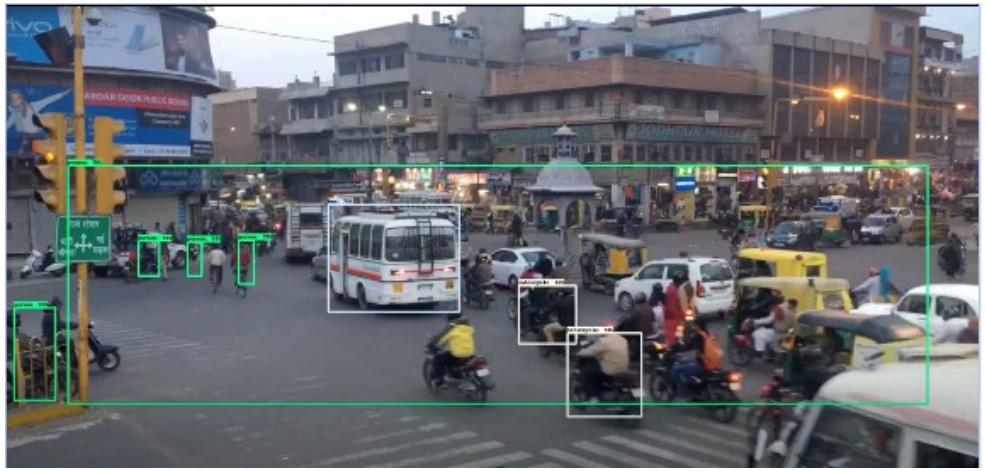

Improve existing system via AI/ML-driven control to manage/configure application modules that can be deployed on smart lampposts and/or to datacenters.

- **Smart Agriculture**

AUGMENTA
PRECISION AGRICULTURE - REDEFINED

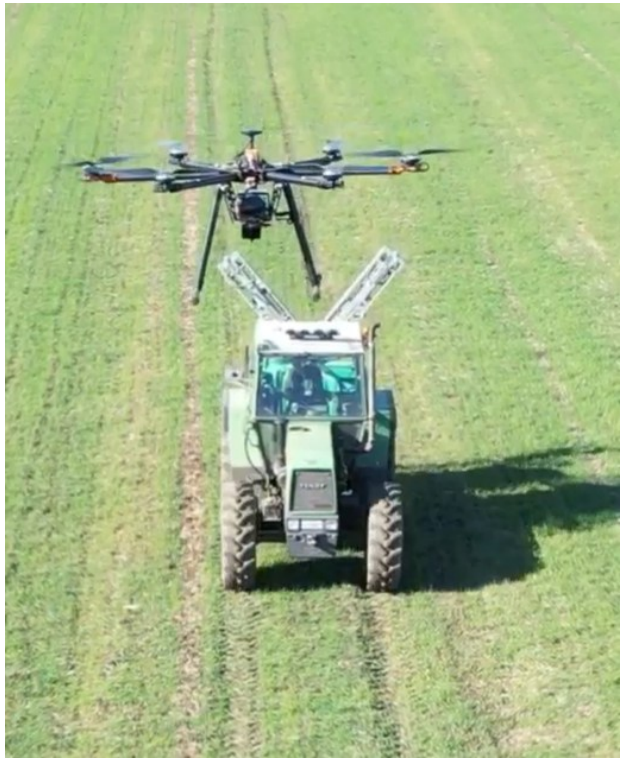

Improve existing system via AI/ML-driven control to jointly manage / configure the image processing application pipeline on a tractor and a companion drone.

ML for system management

Success stories

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

- **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
- Migration counts & SLO violations
  - **Decrease of 41.33% / 34.98%** respectively, compared with SoTA conventional (non-ML) policy

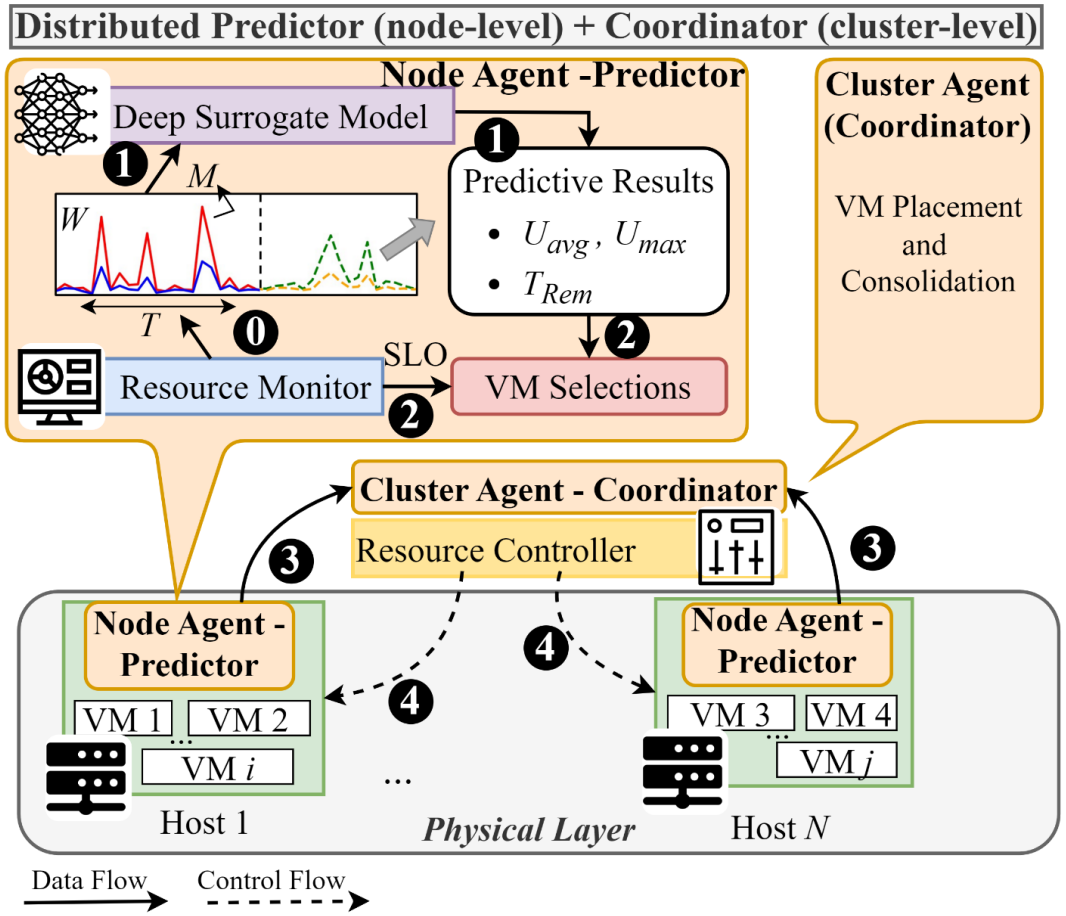

Do we have to train from scratch for each target cluster?

- Trained DRL agents for 4 **different cluster sizes** (4, 8, 16, 32 nodes)
- **Transferred agents between clusters** and compared with training from scratch
- Pre-trained agents consistently outperformed agents trained from scratch when comparing performance over the same training duration
  - **Performance improved by up to 54%**
  - In some cases, pre-trained agents outperformed the converged from-scratch performance in **less than 25% of the training time**

Can ML improve security at the edge?

- **Detect unauthorized / malicious end-devices**
  - Using multiparametric Physical-Level Authentication (PLA), on top of cryptographic authentication solutions
  - Based on unique physical characteristics of the communication subsystem.
- **95% - 98%** malicious devices **detection rate**
  - Even with large number and diverse characteristics of malicious devices…
  - … and for varying SNR levels
- Mean **inference time: 3.75ms**
- **Energy** consumption: **< 25.5mJ**

How to use MLSysOps?

Open-Source Release

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

- Open-source components of the MLSysOps framework (6/2025)
  - Orchestrators
  - Runtimes (sandboxed & generic)
  - Agent templates (continuum, cluster, node)
  - Policy
  - ML Connector
- New release (with additional functionality) expected end-of-year

Open-Source Release Webinar

https://www.youtube.com/@mlsysopsproject

- Hands-on session to introduce the MLSysOps open-source framework (18/6/2025) – **recording available online**
- What will you learn?
  - Why we built the MLSysOps framework and the challenges it addresses
  - How to set up a testbed from scratch using our provided scripts
  - Step-by-step system deployment and execution of a real-world example
  - A demo of our Policy API and a sneak peek at ML integration
  - A look ahead: How you can get involved

Hackathon

https://mlsysops.eu/hackathon/

- Participants will have the opportunity to:
  - Tackle real-world problems focused on the integration of ML with system management
  - Collaborate with researchers, developers, and industry professionals
  - Learn from expert mentors and discover hands-on approaches
  - Connect with an international community of like-minded people.
- 60 registered contestants

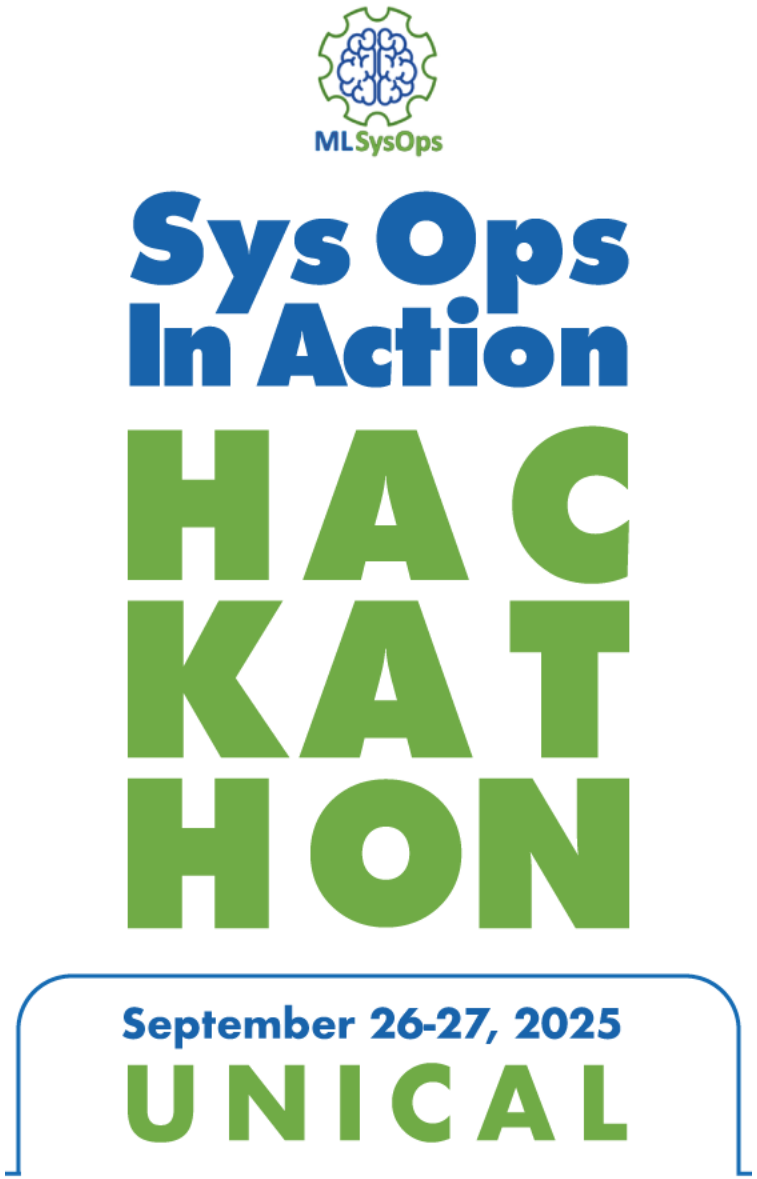

Where to meet us?

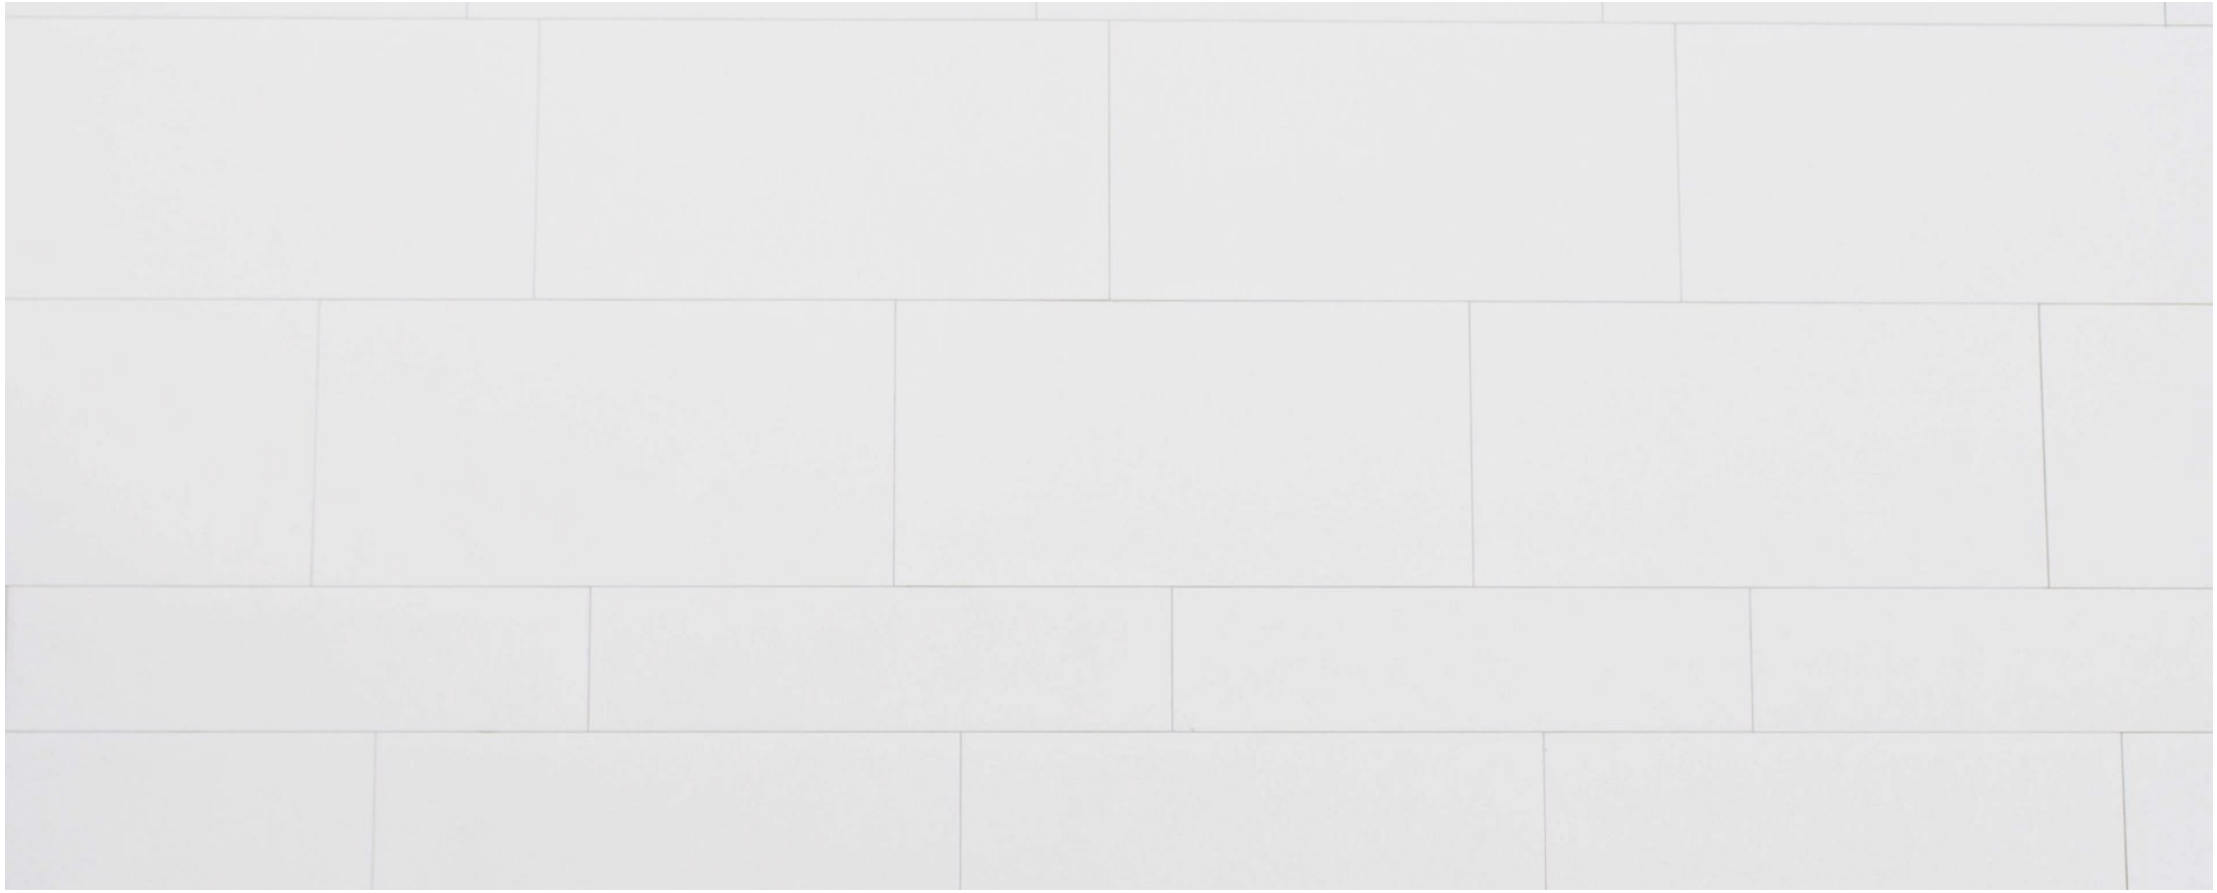

HiPEAC 2026 Workshop

https://ml4ecs.e-ce.uth.gr/

- Workshop organization (ML4ECS - HiPEAC 2026)
  - Edgeless
  - CODECO
  - MLSysOps

Thank you!

https://mlsysops.eu

mlsysops-eu/mlsysops-framework

@mlsysops

mlsysops

@mlsysopsproject

Newsletter https://mlsysops.eu/communication/