Flexible, Scalable, Secure, Decentralized

a MetaOS spanning IoT - Edge - Cloud Continuum

Fulvio Risso
Technical Coordinator FLUIDOS

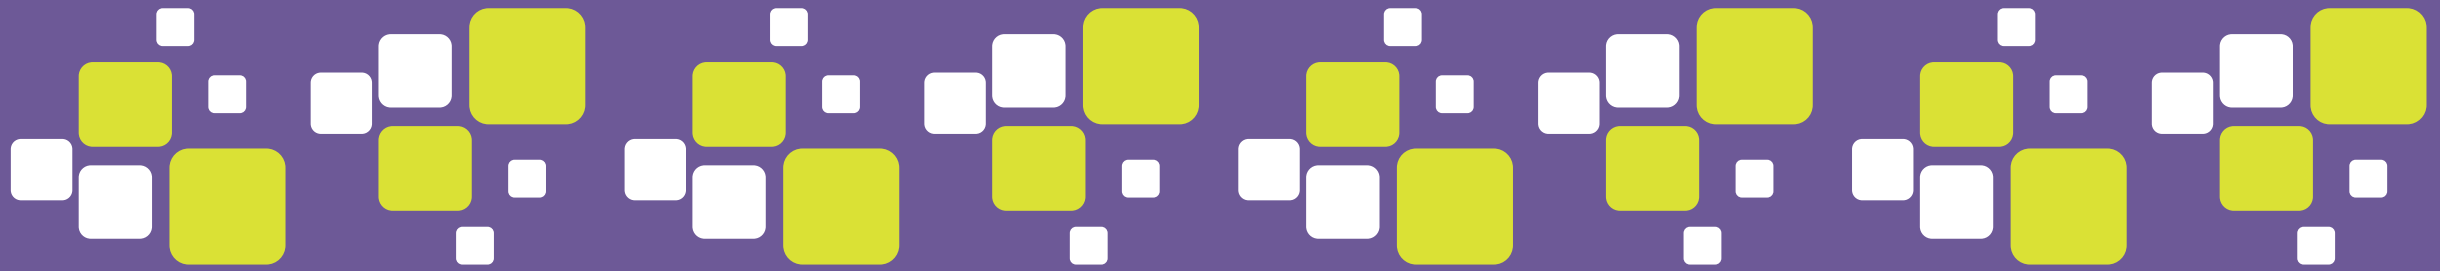

FLUIDOS : A Dynamic and Trustable Computing Continuum

"FLUIDOS (Flexible, scaLable, secUre, and decentralIseD Operating System) aims to leverage the enormous, unused processing capacity at the edge, scattered across heterogeneous edge devices, servers and on-prem datacenters, that struggle to integrate with each other and to coherently form a seamless computing continuum."

FLUIDOS Objectives:

1. Unify edge and cloud computing through decentralized, autonomous resource integration.
2. Shift computing gravity beyond data centers, fostering cross-provider community computing.
3. Orchestrate services and applications across devices and domains with energy-efficient AI.
4. Implement a Zero-Trust security approach for authenticated, authorized access to dispersed resources.
5. Cultivate a multi-stakeholder edge services market, promoting European digital autonomy.

FLUIDOS

FLUIDOS – Why and how

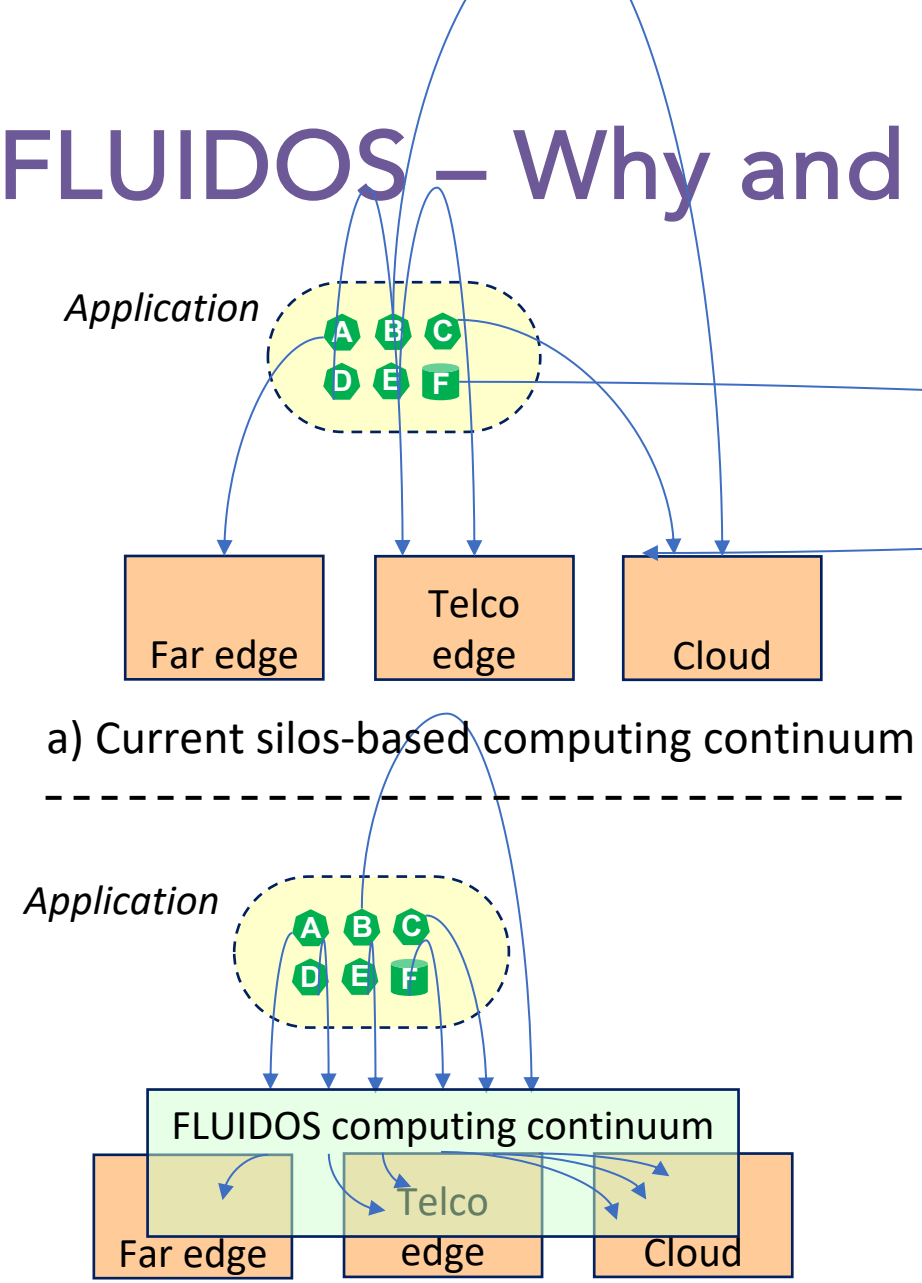

The computing continuum envisioned by FLUIDOS (a.k.a., Liquid Computing) is not simply the capability to deploy services in multiple sites datacenter clusters or devices.

It **defines a virtual space**, spanning across multiple technological domains and, potentially, across multiple administrative boundaries **in which the three properties hold: Deployment transparency, Communication transparency, and Resource availability transparency**.

FLUIDOS – Technologies & Components

Technology

A FLUIDOS node builds on top of **Kubernetes**, which takes care of abstracting the underlying (physical) resources and capabilities in a uniform way and **Liqo.io**, which provides the fundamental primitives for the computing continuum.

FLUIDOS **enriches** the above software with new **control logic**, which handles node-to-node interactions, as well as it enables the specification of advanced policies and intents (e.g., to constrain application execution), which are not supported by the vanilla Kubernetes orchestrator.

Main Components

- **REAR (REsource Advertisement and Reservation) protocol**: it enables different actors such as cloud providers and customers to **advertise**, **reserve** (and then **consume**) *resources* (CPU, RAM; etc.), and *services* (e.g., a database as a service).
- **Available Resources database**: keeps up-to-date information about resources and services available either locally or acquired from remote nodes, following the resource negotiation and acquisition process.
- **Node Orchestrator**: orchestrates service requests, coordinate all the interactions with local components (e.g., local scheduler) and remote nodes (e.g., to set up the computing / network / storage / service fabrics), and make sure that the service behaves as expected.
- Additional modules such as **cluster discovery**, **security manager**, etc.

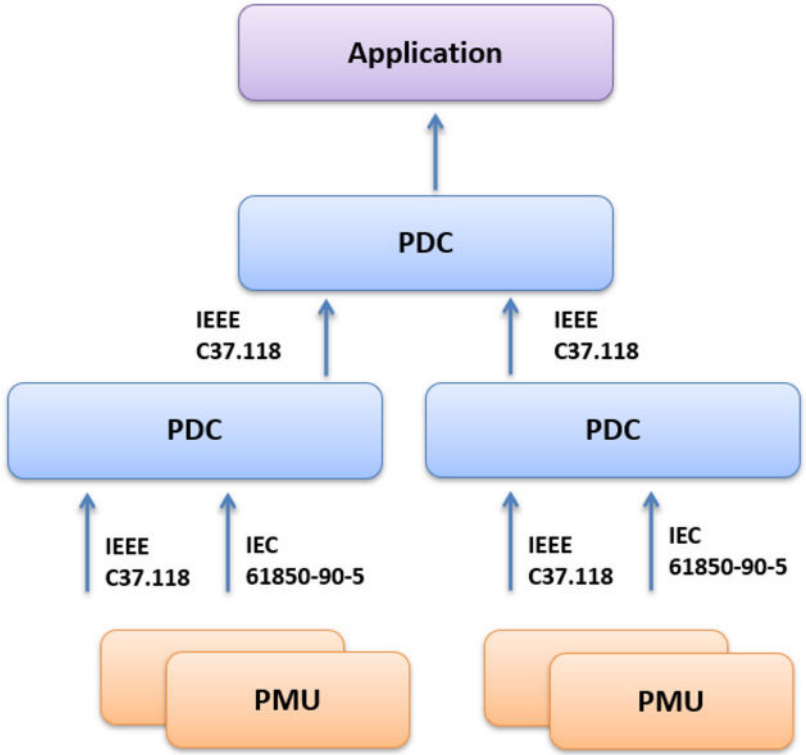

A Phasor Data Concentrator (PDC) is a function responsible for the collection, aggregation and time synchronization of data from Phasor Measurement Unit (PMU) or other PDCs and their transmission to other PDCs or to a target application.

Use Case I : Intelligent Power Grid

- ***Methodology***: Integration of Cloud-Edge processing with electric grids, incorporation of sensors in the continuum, network analytics-based issue detection.
- ***Objective*:** Thanks to FLUIDOS orchestration, we'll have seamless phasor data concentration and grid state computation in presence of multiple ICT outages reducing the risk of electrical downtime.

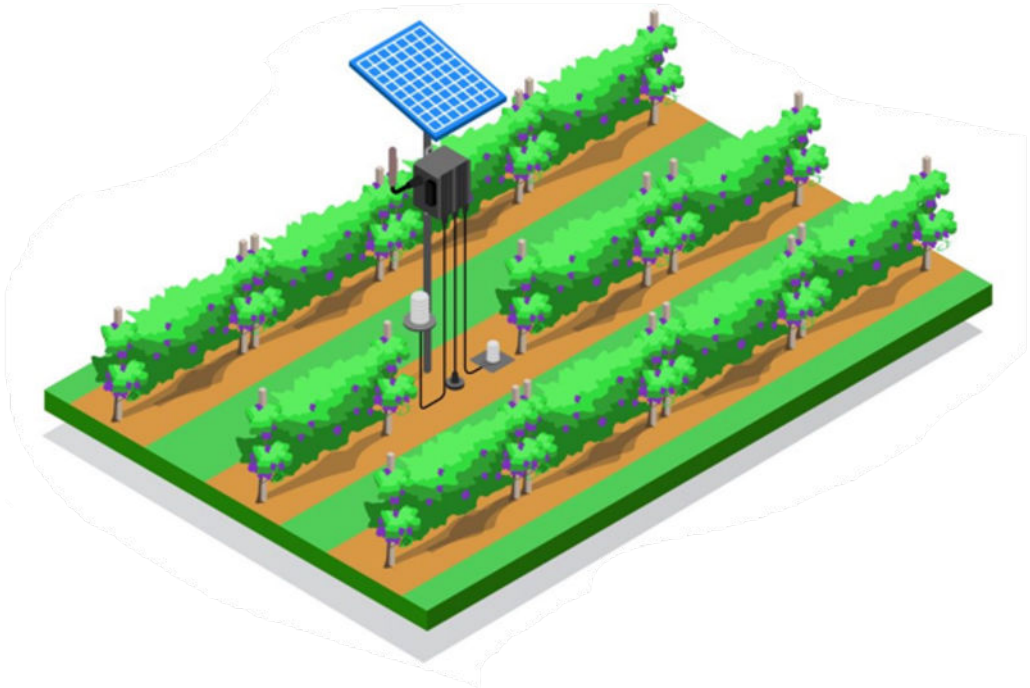

Use Case II : Smart Viticulture

- ***Methodology***: Enabling smart farming for viticulturalist using Terraview's climate SaaS platform bringing data from multiple sources with proprietary AI/ML pipelines to help create intelligence for the practitioners on the ground, to make better decisions and make better wines in a climate-sustainable way.
- ***Objective***: Leverage FLUIDOS to simplify the deployment and the process for critical decisions on the ground using intelligence from various sources, automatized data and analytics with connected reports, automatic alert response system, increased efficiency of workforce, improved health of crops and reduced loss of vines, decreased use of chemicals.

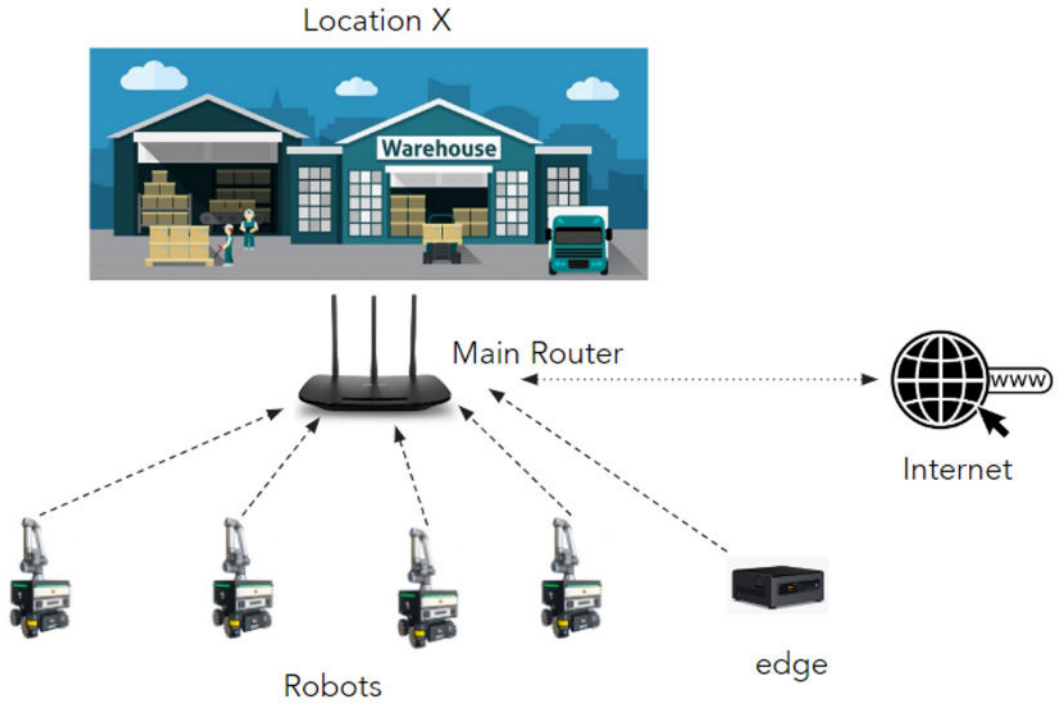

Use Case III : Robotic Logistics

- ***Methodology***: Integration of Cloud-Edge processing with factory robots, implementation of machine learning for coordination and orchestration purposes, optimization of object recognition capabilities in small robots.
- ***Objective***: Integrate FLUIDOS to increase the individual and collective energy efficiency, capacity-based distribution of energy and computing resources, without draining robot batteries and reduce deployment and hardware costs. Possibly, predictive strategies on future energy demand based on past loads and strategies.

FLUIDOS OPEN CALLS

- In FLUIDOS, Open Calls play a crucial role in community building, gathering requirements, and validating the platform. FLUIDOS will provide financial support through two grant types:

1. Technology Extension Grant (TEG)
   a. December 2023 – February 2024

2. Use Case Grant (UCG)
   a. August 2024 – October 2024

1) Technology Extension Grant (TEG): For individual applicants (i.e., no consortia), with a maximum award of €75,000. TEG grants are available in the first Open Call (M16) and are designed for 3rd party projects enhancing FLUIDOS with new features, tested in novel use cases.
   a. TEG will be awarded to 5 applicants

1) Use Case Grant (UCG): Available for both individuals and small consortia (up to 3 entities), with a maximum award of €120,000. It is for 3rd party projects creating demonstrators to validate FLUIDOS architecture and software.
   a. UCG will be awarded to 10 applicants:
      - 5 applicants in the first Open Call
      - 5 in the second Open Call

FLUIDOS – EVALUATION TIMELINE

Evaluation Process

1. The proposal evaluation process aims to be efficient, taking about a month.
2. The evaluation process involves external and project-based technical partner evaluations
3. There will be evaluators' open call launching in the next week.
4. The selected evaluators will sign non-disclosure and impartiality agreements.
5. Evaluation templates are provided after an online introduction meeting and agreement signing.
6. Evaluators review applications with identifying information removed within three working days.
7. The final selection of successful applicants, 10 for the first Open Call and 5 for the second, is made in a consensus meeting within three weeks of the call's closure.

https://fluidos.eu

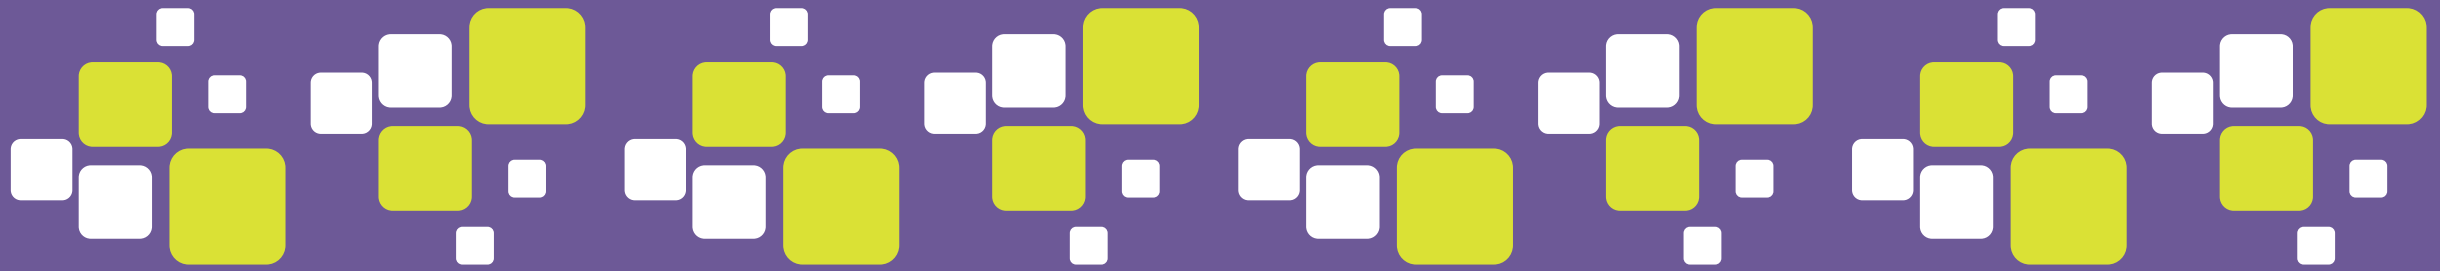

FLUIDOS - Architecture

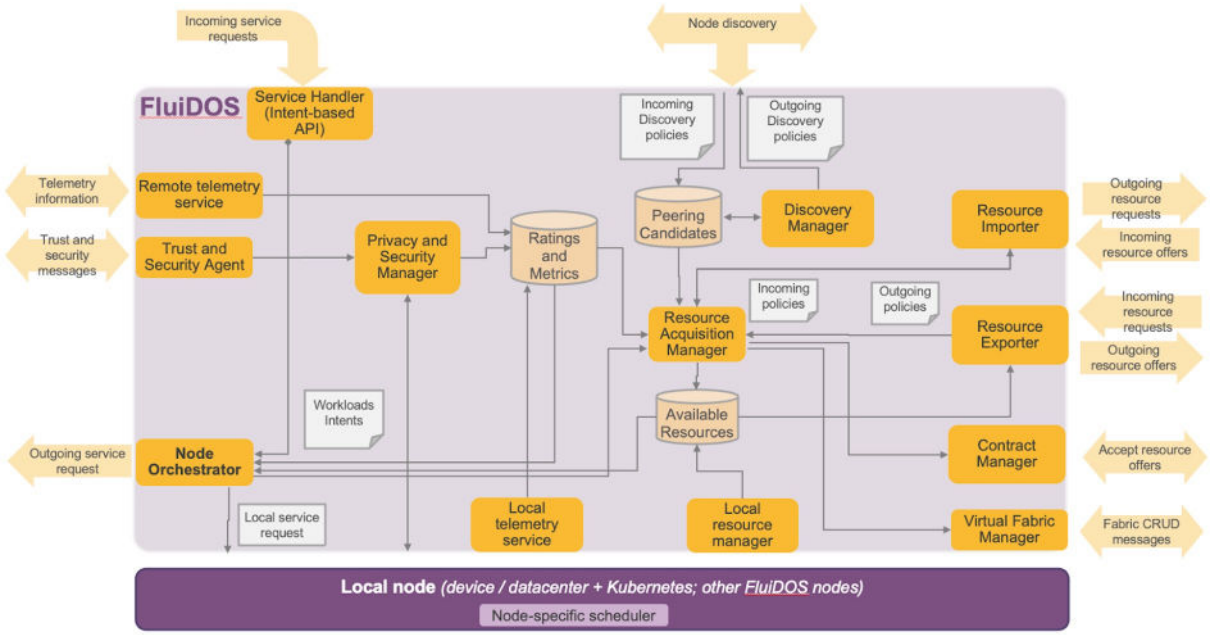

The *discovery manager* is the module responsible for the discovery of other FLUIDOS nodes, producing as output a local database of feasible *peering candidates*.

Node characteristics are exposed at a high level (e.g., through generic key/value labels); they enable both an initial policy driven filtering (i.e., excluding undesired nodes from the list of *peering candidates*) and a-priori filtering and ranking during the resource acquisition phase, while leaving that phase to deal with the negotiation in terms of quantities, cost, and more detailed aspects.