

# Enabling Reproducible Analysis of Complex Application Workflows on the Edge-to-Cloud Continuum

**PhD candidate**

*Daniel Rosendo*

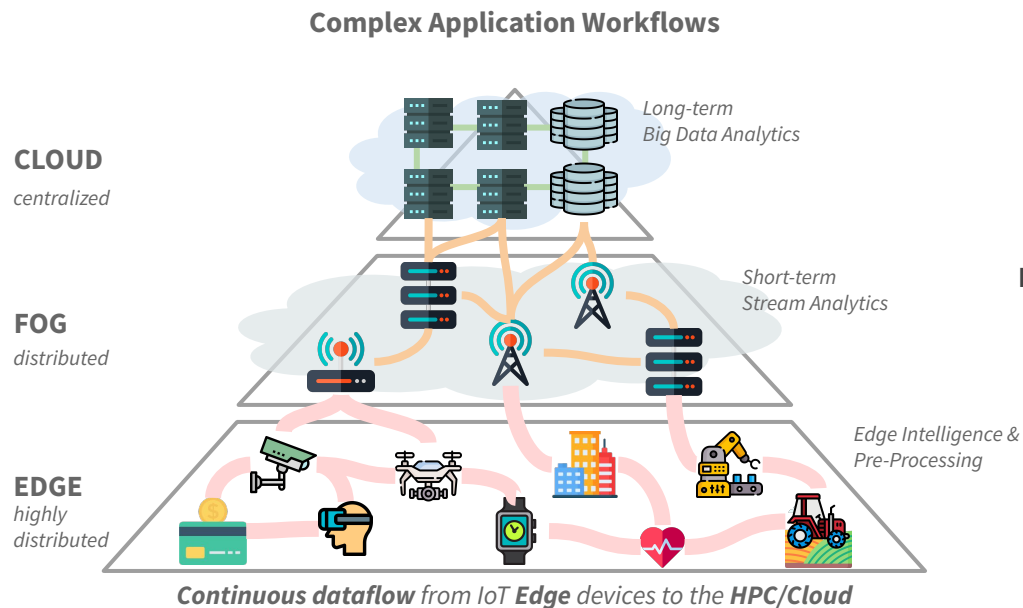
**Advisors**

*Gabriel Antoniu*

*Alexandru Costan*

*Patrick Valduriez*

# The Computing Continuum



## Research Challenges & Opportunities



**Understanding Performance of**  
Application Workflows on the  
Edge-to-Cloud Continuum

**Optimizing Performance of**  
Edge-to-Cloud Application Workflows

**Supporting Reproducible Analysis of**  
Complex Edge-to-Cloud Workflows

# Research Challenge 1: Understanding Performance

## Research Challenges & Opportunities

*Which system **parameters** and infrastructure **configurations** impact performance and how?*

*How to **systematically** perform large-scale **experiments** to enable the **reproducibility** of the **results**?*



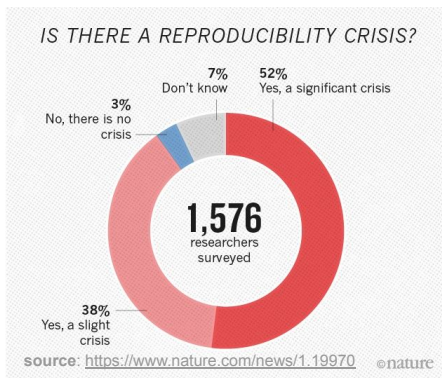
## Evaluation & Validation



Representative setup to understand performance



Repeatable, Replicable & Reproducible experiments

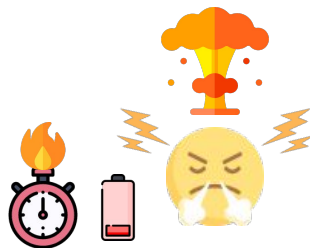


***“+70% failed to reproduce another scientist's experiments”***

***“+50% failed to reproduce their own experiments”***



## Technical Challenges



Install & configure services

Configure network

Map application parts with underlying infrastructure

Define the execution workflow

Establish procedures for reproducibility



## Simple testbed setups



Many abstractions

Scalability issues

Hard to reproduce

# What Would Be an Ideal Solution?

## Research Challenges & Opportunities

*Which system **parameters** and infrastructure **configurations** impact performance and how?*

*How to **systematically** perform large-scale **experiments** to enable the **reproducibility** of the **results**?*



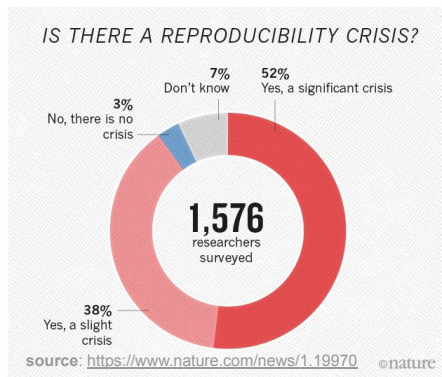
## Evaluation & Validation



Representative setup to understand performance



Repeatable, Replicable & Reproducible experiments



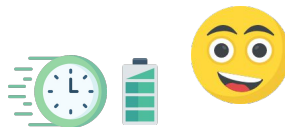
***“+70% failed to reproduce another scientist's experiments”***

***“+50% failed to reproduce their own experiments”***



## Technical Challenges

**Focus on high-level aspects**



Install & configure services

Configure network

Model with **Abstract complexities**

Define the execution workflow

Establish procedures for reproducibility



## Large-scale testbed setups

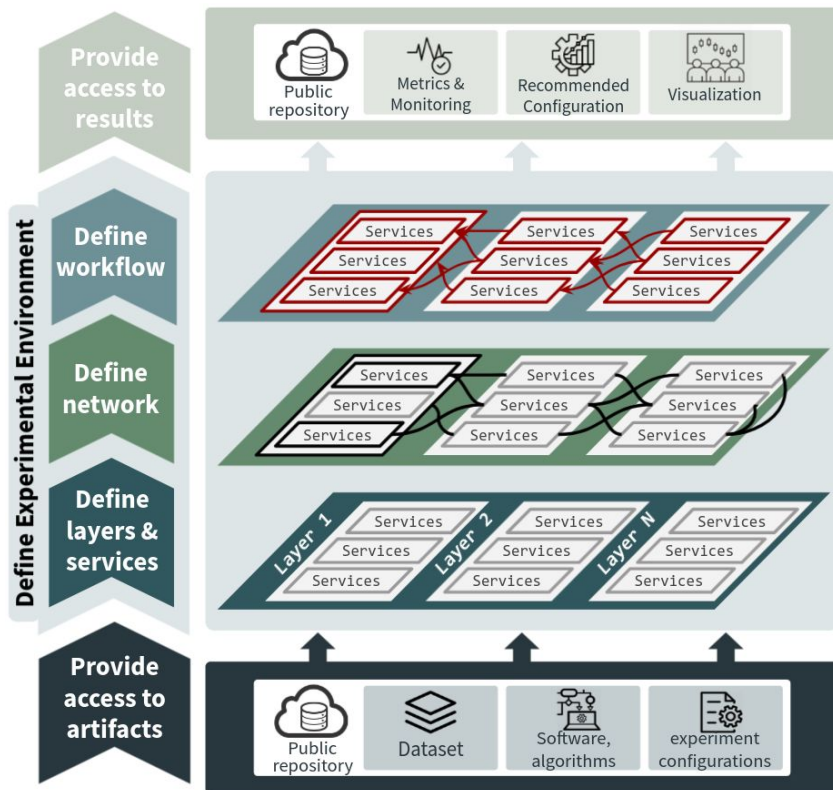


Representative setup

Easy to scale

Supports reproducible experiments

# What is E2Clab?



**IEEE Cluster 2020**

<https://hal.archives-ouvertes.fr/hal-02916032>

## Methodology



### Reproducible Experiments

*Repeatability, Replicability & Reproducibility*



### Mapping

*Application parts & physical testbed*



### Variation & Scaling

*Experiment variation and transparent scaling*



### Network Emulation

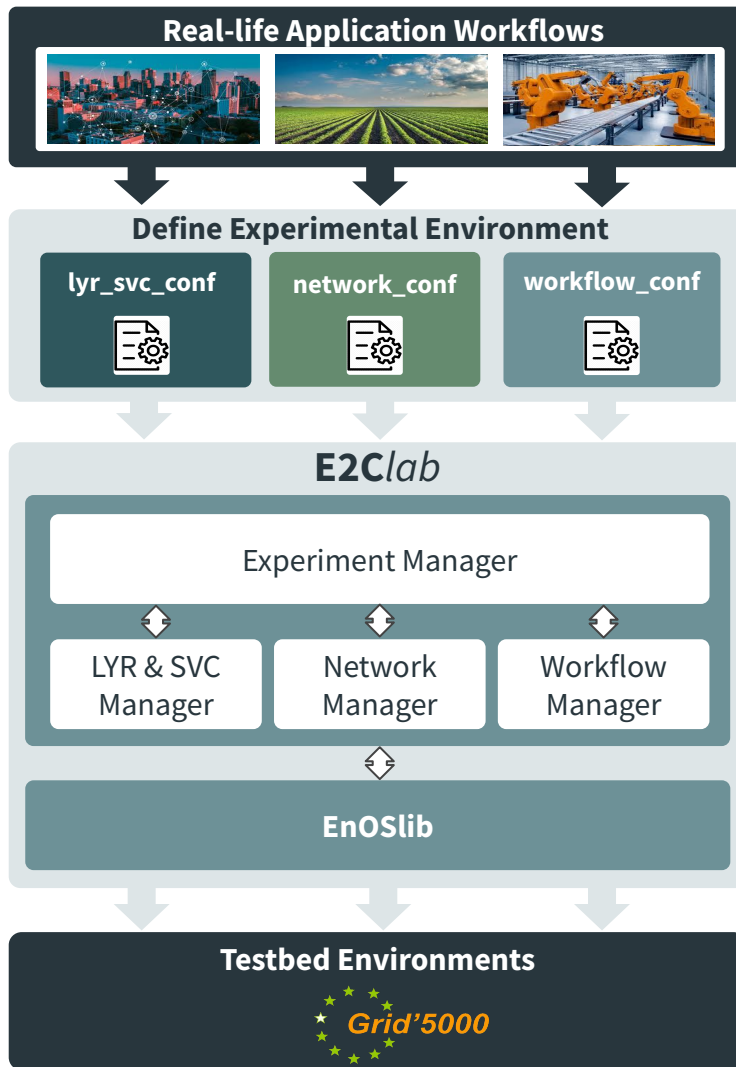
*Edge-to-Cloud communication constraints*



### Experiment Management

*Deployment, Execution & Monitoring*

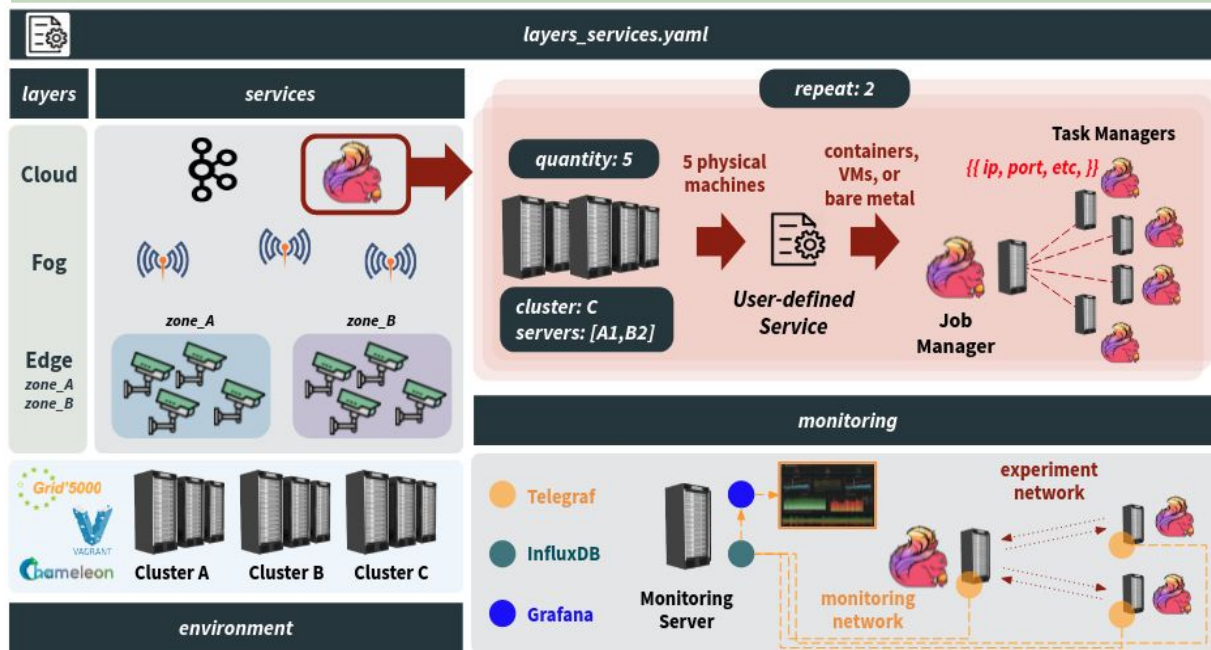
# What is **E2Clab**?



# Defining the experimental environment: *Layers and Services*

```
1 environment:
2   name: g5k
3   job_name: example
4   env_name: "debian10-x64-big"
5   walltime: "01:30:00"
6 monitoring:
7   type: tig
8   network: private
9 layers:
10 - name: cloud
11   services:
12     - name: Flink
13       cluster: chiclet
14       quantity: 5
15       repeat: 2
16       roles: [monitoring]
17     - name: Kafka
18       servers: ["chifflet-8.lille.grid5000.fr"]
19       roles: [monitoring]
20 - name: fog
21   services:
22     - name: Mosquitto
23       quantity: 3
24       roles: [monitoring]
25 - name: edge.zone_A
26   services:
27     - name: Cameras
28       quantity: 4
29 - name: edge.zone_B
30   services:
31     - name: Cameras
32       quantity: 4
```

## Selecting, reserving, monitoring, and configuring resources





# Defining the experimental environment: *User-defined Service*

*“Our methodology can be generalized to other applications!”*

## Reusable services



ClientKerA.py

Default.py

Flink.py

Horovod.py

Kafka.py

KerA.py

KerAF.py

KerFlink.py

Mosquitto.py

PlantNetEngine.py

```
1 from e2clab.services import Service
2
3 class Flink(Service):
4
5     def deploy(self):
6         # service logic
7         self.deploy_docker(registry_opts=registry_opts)
8
9         JM_PORT = "8081"
10        job_manager_service = "job_manager"
11        task_manager_service = "task_manager"
12
13        roles_job_manager = Roles({job_manager_service: [self.hosts[0]]})
14        roles_task_manager = Roles({task_manager_service: self.hosts[1:len(self.hosts)]})
15
16        JOB_MANAGER_ADDRESS = roles_job_manager[job_manager_service][0].address
17        self.env.update({"JOB_MANAGER_RPC_ADDRESS": JOB_MANAGER_ADDRESS})
18
19        with actions(roles=roles_job_manager) as p:
20            p.docker_container(name=job_manager_service,
21                              image="flink:1.9-scala_2.11",
22                              restart="yes",
23                              restart_policy="always",
24                              network_mode="host",
25                              env=self.env,
26                              command="jobmanager")
27
28        with actions(roles=roles_task_manager) as p:
29            p.docker_container(name=task_manager_service,
30                              image="flink:1.9-scala_2.11",
31                              restart="yes",
32                              restart_policy="always",
33                              network_mode="host",
34                              env=self.env,
35                              command="taskmanager")
36
37        # register service
38        self.register_service(_roles=roles_job_manager, service_port=JM_PORT, sub_service=job_manager_service)
39        self.register_service(_roles=roles_task_manager, sub_service=task_manager_service)
40
41        return self.service_extra_info, self.service_roles
```

Big Data  
frameworks

AI  
frameworks

User's  
applications



*among others!!!*



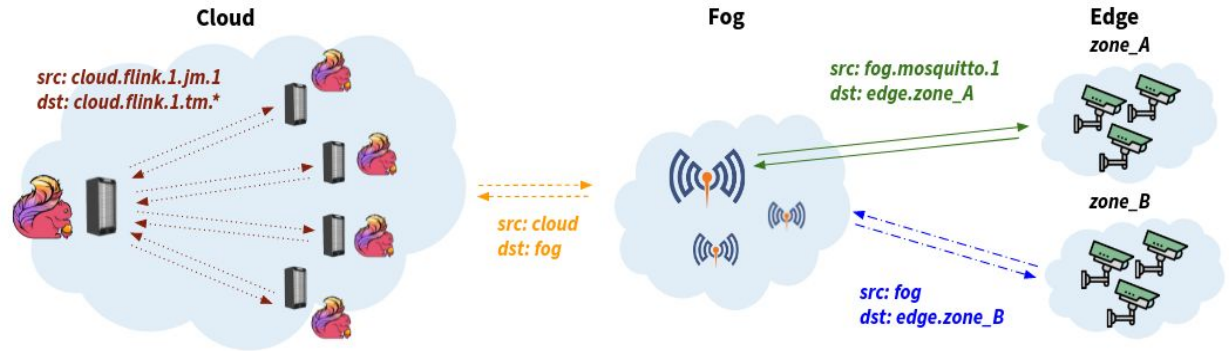
# Defining the experimental environment: *Network*

```
1 networks:
2 - src: cloud
3   dst: fog
4   delay: "2ms"
5   rate: "10gbit"
6   loss: 2
7 - src: fog.mosquitto.1
8   dst: edge.zone_A
9   delay: "2ms"
10  rate: "1gbit"
11  loss: 5
12 - src: fog
13   dst: edge.zone_B
14   delay: "20ms"
15   rate: "150mbit"
16   loss: 7
```

## Defining network constraints



network.yaml



# Defining the experimental environment: *Workflow*

```
1 - hosts: edge.zone_A.*
2   depends_on:
3     service_selector: "fog.mosquitto.*"
4     grouping: "round_robin"
5     prefix: "mosquitto"
6   prepare:
7     - copy:
8       src: "{{ working_dir }}/libs/"
9       dest: "/opt/lib/"
10  launch:
11    - shell: "java my-app.jar tcp://{{ mosquitto.url }}
12              /opt/metrics/throughput"
13  finalize:
14    - fetch:
15      src: "/opt/metrics/throughput"
16      dest: "{{ working_dir }}/experiment-results/producers/"
17 - hosts: ...
18 ...
19 - hosts: ...
```

ANSIBLE

ANSIBLE

ANSIBLE

## Defining the execution flow and relationships

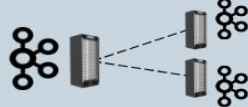


workflow.yaml

hosts

cloud.\*

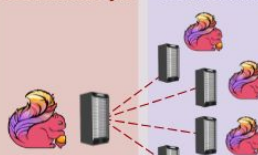
cloud.kafka.\*



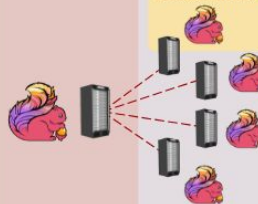
cloud.flink.\*

cloud.flink.\*.jm

cloud.flink.1.tm.\*



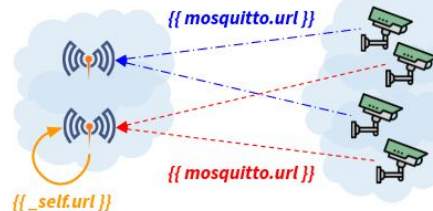
cloud.flink.2.tm.1



depends\_on

Fog

Edge zone\_A



prepare



Dataset



Software, algorithms



experiment configs.

launch



finalize



Metrics & Monitoring

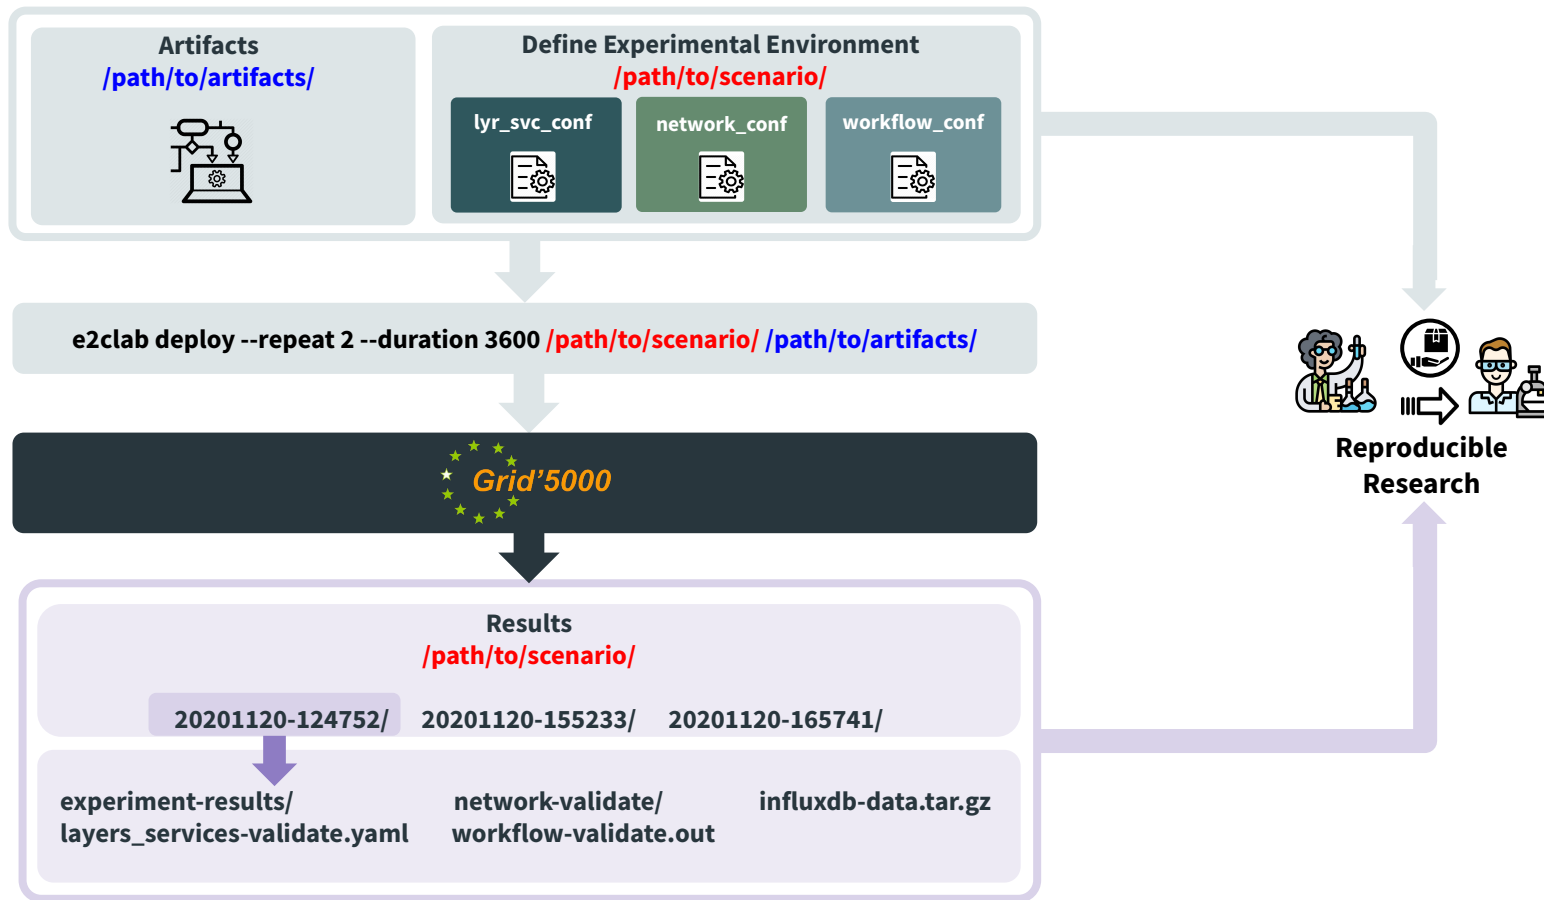


Optimal Configuration



Visualization

# Experimental cycle in **E2Clab**



# Research Challenge 2: Optimizing Performance

*How to configure the system components to minimize the processing latency?*

## **Metrics**

energy, cost, resource usage & latency

## **Objective**

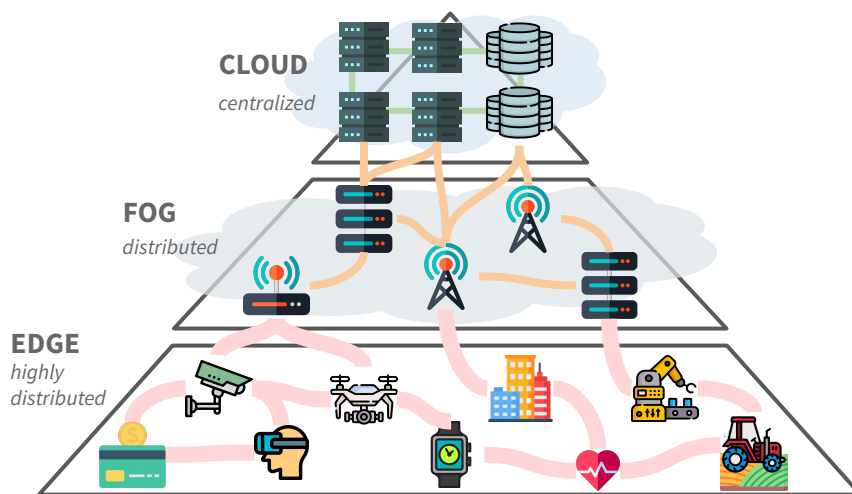
minimize processing time

## **Constraints**

energy consumption limits & equipment cost < budget

## **Variables**

config. params. of Edge, Fog, and Cloud systems



*Where should application parts be executed to minimize communication costs and end-to-end latency?*

## **Metrics**

costs, latency, & resource usage

## **Objectives**

minimize communication costs & minimize end-to-end latency

## **Constraints**

budget & hardware resource limits

## **Variables**

Edge-to-Cloud network communication links (bandwidth and delay) & number of Fog nodes per Edge device

# Problem Statement



*The performance optimization of application workflows  
on highly heterogeneous resources is challenging!*



## **Heterogeneous constraints**

computing resources;  
network communication;  
application requirements;



## **Search space complexity**

myriad of combination of  
possibilities;  
multi-objective problems;  
NP-hard complex;



## **Selecting the most accurate optimization technique**

multiple methods that perform  
differently;  
hyperparameter search;

# Our Optimization Methodology



IEEE Cluster 2021

<https://hal.archives-ouvertes.fr/hal-03310540>

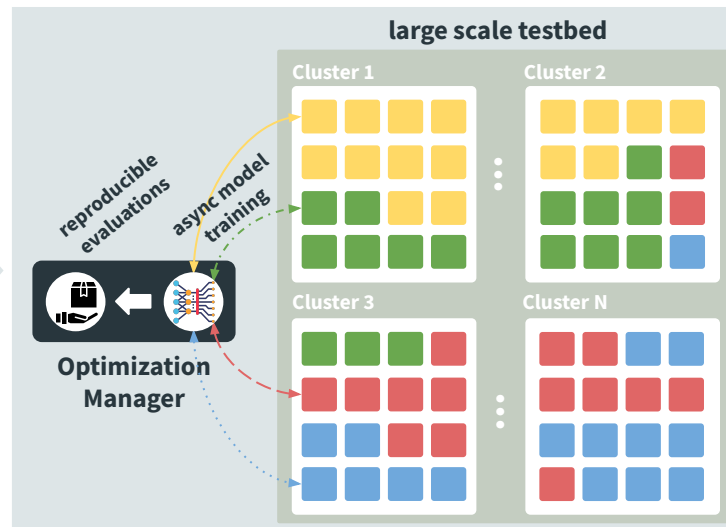
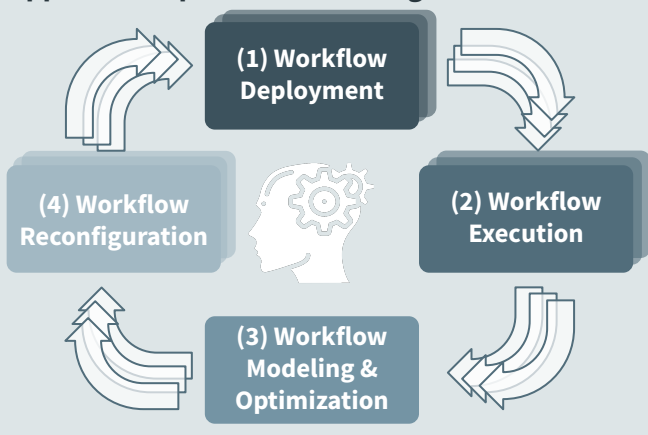
Phase I

Define Optimization Problem

$$\begin{aligned} \min/\max_x \quad & f_m(x), \quad m = 1, 2, \dots, M \\ \text{subject to} \quad & g_j(x) \leq 0, \quad j = 1, 2, \dots, J \quad \text{Inequality constraints.} \\ & h_k(x) = 0, \quad k = 1, 2, \dots, K \quad \text{Equality constraints.} \\ & x_i^L \leq x_i \leq x_i^U, \quad i = 1, 2, \dots, I \quad \text{Bounds on variables.} \end{aligned}$$

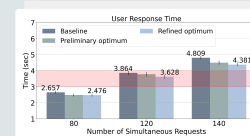
Phase II

**Parallel + Scalable + Reproducible**  
application optimization on large scale testbeds



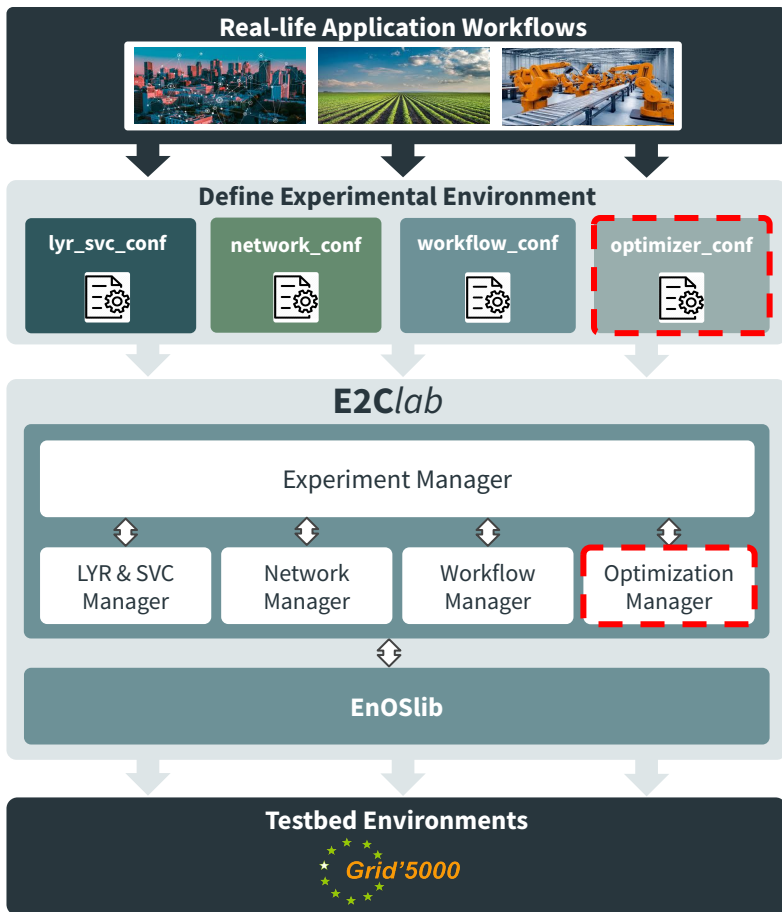
Phase III

Summary of Computations



Reproducible Research

# Implementation in E2Clab



## How to setup an optimization?

### User-defined optimization

```
from e2clab.optimizer import Optimization

class UserDefinedOptimization(Optimization):

    def run(self):
        algo = SkOptSearch(
            optimizer=Optimizer(
                base_estimator='ET',
                n_initial_points=45,
                initial_point_generator="lhs",
                acq_func="gp_hedge"))
        algo = ConcurrencyLimiter(algo,
            max_concurrent=2)
        scheduler = AsyncHyperBandScheduler()
        objective = tune.run(
            self.run_objective,
            metric="user_resp_time",
            mode="min",
            name="plantnet_engine",
            search_alg=algo,
            scheduler=scheduler,
            num_samples=10,
            config={
                "http": tune.randint(20, 60),
                "download": tune.randint(20, 60),
                "simsearch": tune.randint(20, 60),
                "extrac": tune.randint(3, 9)})

    def run_objective(self, _config):
        # create an optimization directory
        self.prepare()
        # deploy the configs on the testbed
        self.launch()
        # backup the optimization computations
        self.finalize()
        # report the metric value to Ray Tune
        tune.report(user_resp_time=user_resp_time)
```



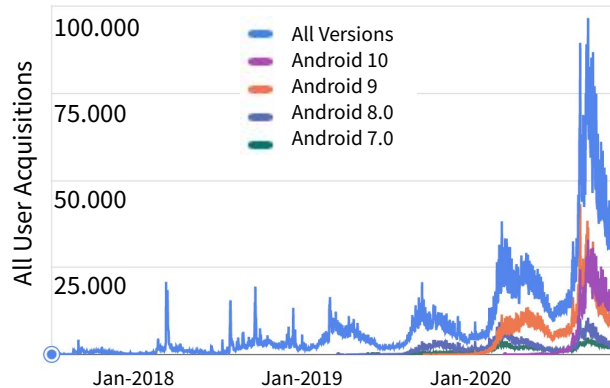
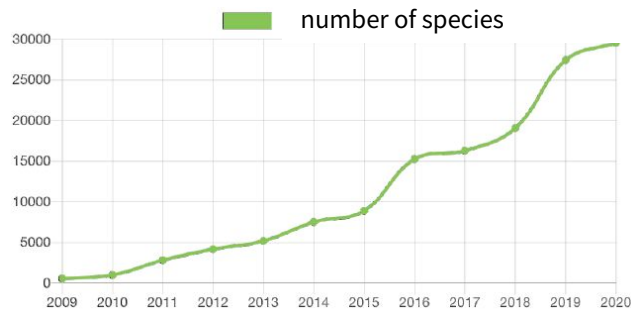
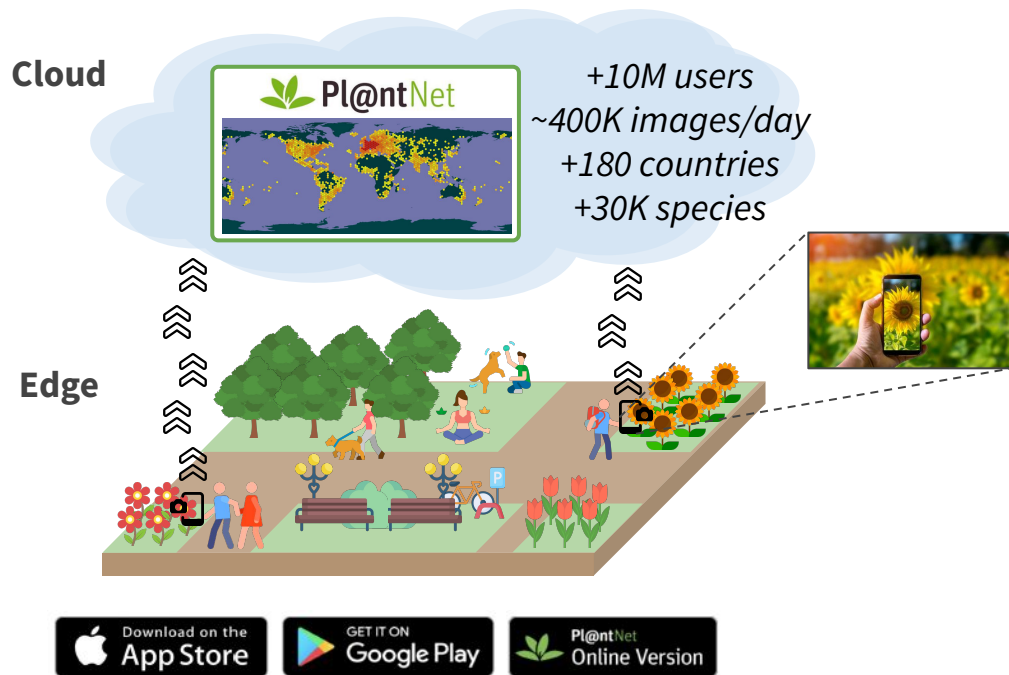
**Supports state-of-the-art Bayesian Optimization libraries**



**among others!!!**



# Validation with a Large-scale Real-life Application: Pl@ntNet



Pl@ntNet has an **exponential grow** of new users **acquisition** every spring (peaks in May-June)



Understanding and Optimizing the Performance of Pl@ntNet

# Validation with a Large-scale Real-life Application: Pl@ntNet

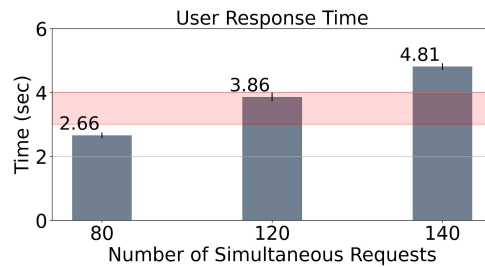
## Pl@ntNet Identification Engine

Thread pool	baseline (# threads)	Description	Hardware
HTTP	40	# simultaneous requests being processed.	CPU
Download	40	# simultaneous images being downloaded.	CPU
Extract	7	# simultaneous inferences in a single GPU.	GPU
Simsearch	40	# simultaneous similarity search.	CPU

## Main performance metric: user response time



“Over **3-4 seconds** more than **60%** of users **abandon** the transaction and may even **delete** the application”



## Research Questions

What is the **software configuration** that **minimizes the user response time**?

How do the **Extraction** and **Similarity Search** thread pool configurations impact the **user response time**?

How does the **number of simultaneous users** accessing the application **impact on the user response time**?



**Large scale experimental evaluations**



**Grid'5000**

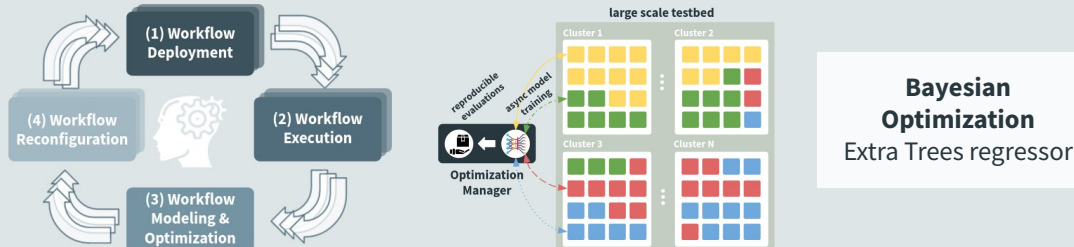
**42 nodes**

# What is the **software configuration** that **minimizes the user response time**?

## Phase I Define Optimization Problem

Find  $(http, download, simsearch, extract)$ , in order to  
Minimize  $UserResponseTime$   
Subject to  $20 \leq (http, download, simsearch) \leq 60$ , Pool Size.  
 $3 \leq (extract) \leq 9$ , Pool Size.

## Phase II Parallel, Scalable, Reproducible Workflow Optimization



## Phase III Summary of Computations

Thread pool	baseline	preliminary optimum
HTTP	40	54
Download	40	54
Extract	7	7
Simsearch	40	53
User response time (sec)	2.65	2.48

### preliminary optimum

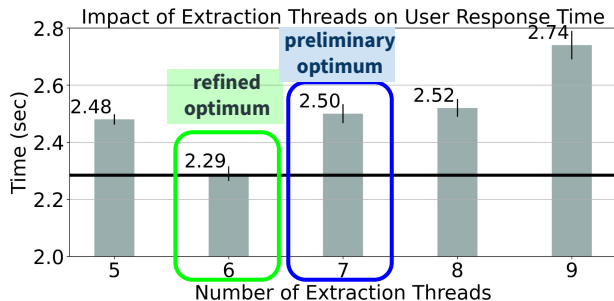
serves 35% more simultaneous users

reduces the user response time by 7%

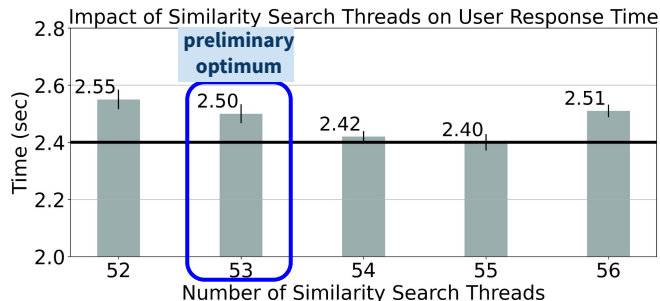
# How do the **Extraction** and **Similarity Search** thread pool configurations impact the **user response time**?

## Sensitivity Analysis: One-at-a-time (OAT) method

### Extraction



### Similarity Search



Thread pool	baseline	preliminary optimum	refined optimum
HTTP	40	54	54
Download	40	54	54
Extract	7	7 [5, 6, 8, 9]	6
Simsearch	40	53 [50, 51, 52, 54, 55, 56]	53
User response time (sec)	2.65 (±0.0914)	2.48 (±0.0912)	2.47 (±0.0826)

preliminary optimum

serves 35% more  
simultaneous users

reduces the user response  
time by 7%



refined optimum

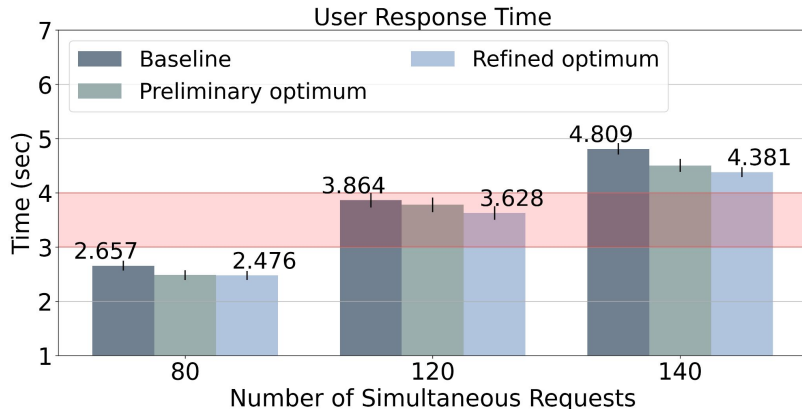
reduces the GPU  
memory usage by 12%

# How does the **number of simultaneous users** accessing the application **impact on the user response time**?

Thread pool	baseline	preliminary optimum	refined optimum
HTTP	40	54	54
Download	40	54	54
Extract	7	7	6
Simsearch	40	53	53

refined optimum
serves 35% more simultaneous users
reduces the user response time by 7%
reduces the GPU memory usage by 12%

“Over **3-4 seconds** more than **60%** of users **abandon** the transaction and may even **delete** the application”



# Main takeaways

**Our methodology** has proved useful for **understanding** and **improving** the **performance** of a **real life application** used at very large-scale.

The **Pl@ntNet** configuration found **using our methodology**  
serves **35% more simultaneous users**  
and **reduces the user response time by 7%**  
compared to the baseline.

# Feedback to G5K staff: **availability of GPU nodes was the main limitation**

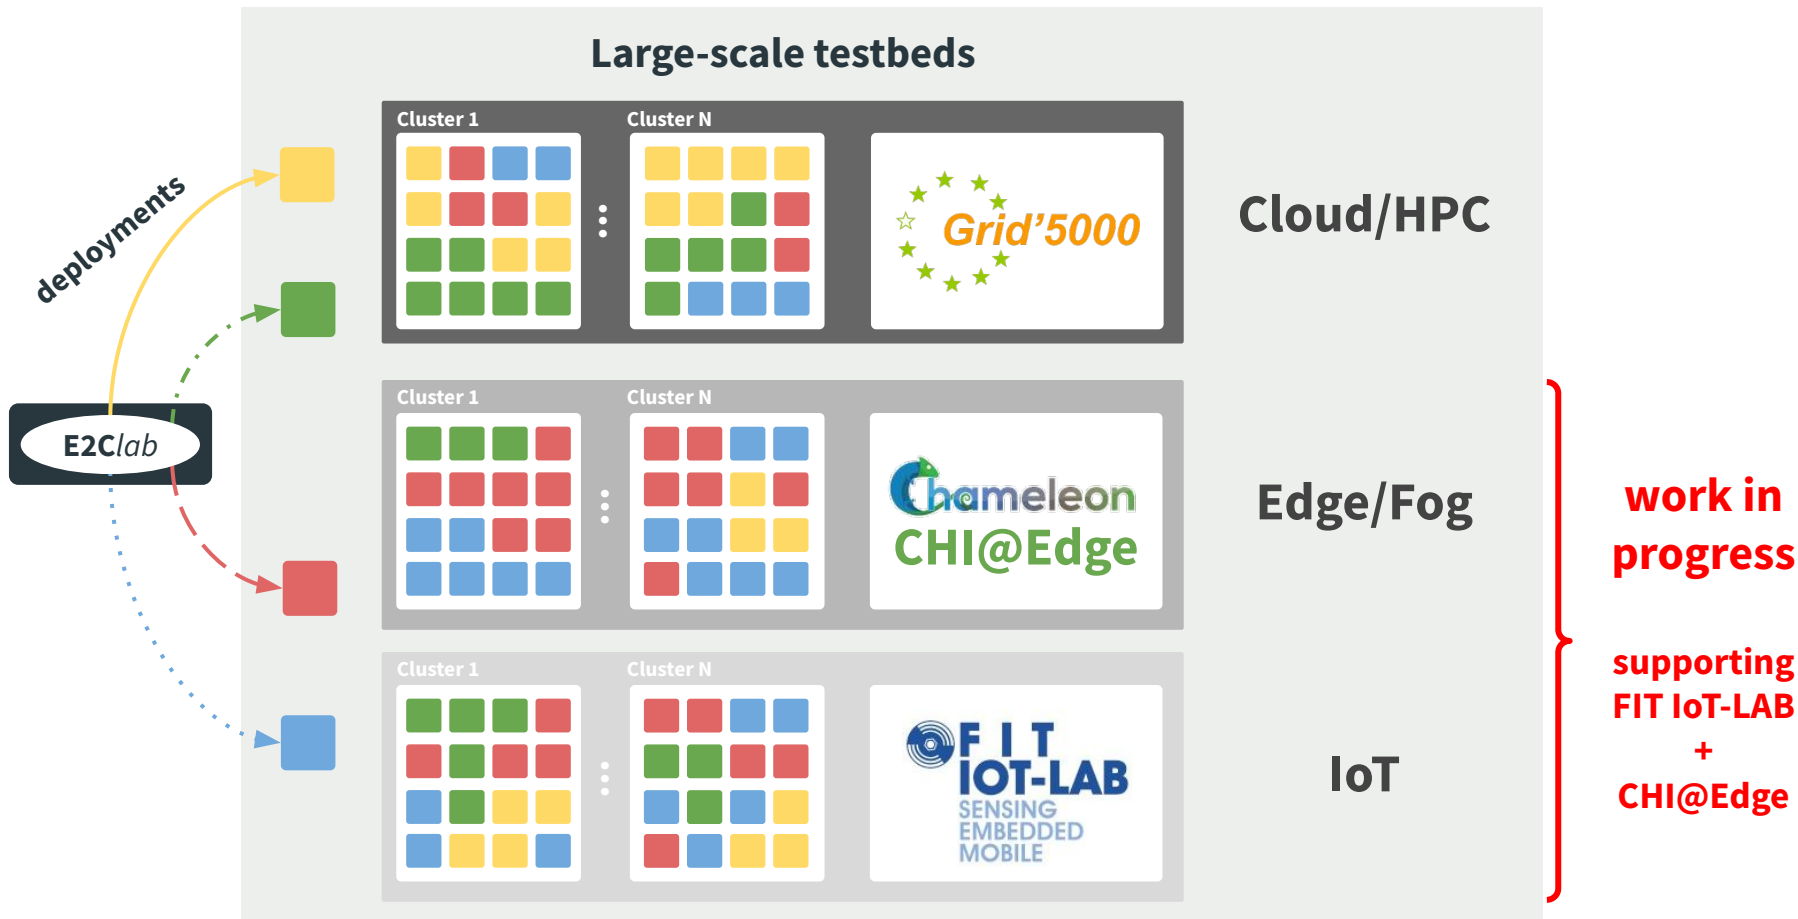
Site	Cluster	# Nodes	GPU	# GPUs	RAM	NIC
Lille	chifflet	8	Nvidia GTX 1080 Ti	2	11GB	eth0/eno1, eth1/eno2
Lille	chifflet	6 2	[1-6] Nvidia Tesla P100 [7-8] Nvidia Tesla V100-PCIE	2 2	16GB 32GB	eth0/ens5f0, eth1/ens5f1
Lyon	gemini	2	Nvidia Tesla V100-SXM2	8	32GB	eth0/enp1s0f0
Lyon	orion	4	Nvidia Tesla M2075	1	5GB	eth0/enp68s0f0
Lyon	neowise	10	AMD Radeon Instinct MI50	8	32GB	eth0/eno1, eth1/eno2
Grenoble	drac	12	Nvidia Tesla P100 (CPU=Power POWER8NVL 1.0)	4	16GB	eth0/enP1p1s0f0
Nancy	grouille	2	Nvidia A100-PCIE	2	40GB	eth2/eno33
Nancy	gruss	4	Nvidia A40	2	45GB	eth2/eno33
Nancy	graffiti	12 1	[1-12] Nvidia RTX 2080 Ti [13] Nvidia Quadro RTX 6000	4 4	11 GB 22 GB	eth2/ens4f0
Nancy	graphique	5	Nvidia GTX 980	2	4GB	eth0/eno1
Nancy	grele	14	Nvidia GTX 1080 Ti	2	11GB	eth0/eno1
Nancy	grimani	6	Nvidia Tesla K40M	2	12GB	eth0/eno1
Nancy	grue	5	Nvidia Tesla T4	4	15GB	eth0/eno1

**Pl@ntNet requires ~27GB of GPU RAM**

Source (October 1st 2020 / December 2nd, 2021): <https://www.grid5000.fr/w/Hardware>



# Ongoing work: from IoT/Edge to Cloud/HPC environments



# Ongoing work: **Provenance capture in E2Clab**

**Provenance capture** for Edge-to-Cloud experiments as a **support to the analysis and debugging** of experiment results

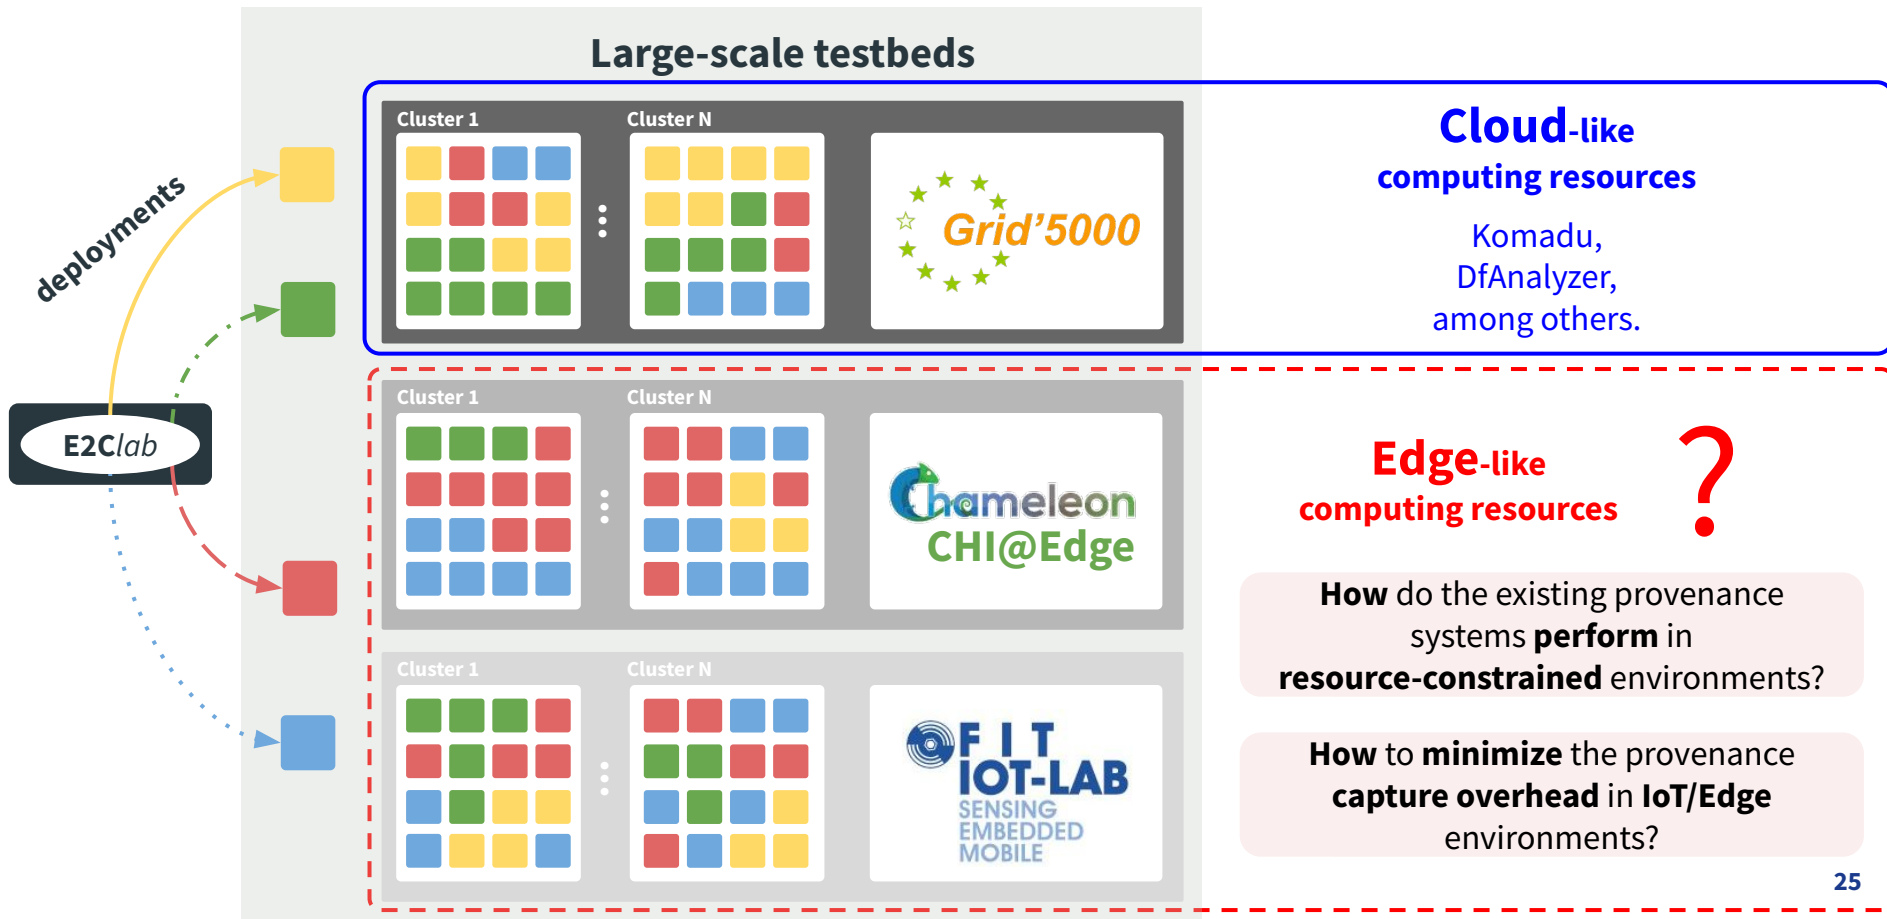
**Goal:** understand how experiment results have been produced

- What **parameters produced** these **results**?
- What **steps** did I **invoke** during **workflow execution**?

**In collaboration with:**

- Marta Mattoso (Federal University of Rio de Janeiro, Brazil)

# Ongoing work: Provenance capture



# Final Considerations



## Reproducible Research



- Access to the experiment **artifacts**, **results**, and **definition of the experimental environment**
  - <https://gitlab.inria.fr/E2Clab/Paper-Artifacts/plantnet>

- **E2Clab** is **open source!** <https://gitlab.inria.fr/E2Clab/e2clab>
- **Documentation:** <https://e2clab.gitlabpages.inria.fr/e2clab/>



## Acknowledgments

- **Inria**
  - Matthieu Simonin, Alexandru Costan, Gabriel Antoniu, Patrick Valduriez
- **Hasso-Plattner-Institut**
  - Pedro Silva
- **Pl@ntNet team**
  - Jean-Christophe Lombardo, Alexis Joly
- **Argonne National Laboratory**
  - Romain Egele, Jaehoon Koo, Prasanna Balaprakash, Orcun Yildiz
- **PhD funded by:** HPC-BigData Inria Project LAB

# Thank you!