IOT-LAB and GRID'5000

Simon Delamare (LIP/CNRS)

Guillaume Schreiner (ICube/CNRS)

RSD AUTUMN SCHOOL
2021/10/4

Plan

- Scientific context
- SILECS testbeds overview
  - IoT-LAB
  - Grid’5000
- About reproducible experiments
- Conclusion

Scientific context

Scientific issues

- Network protocols
- End-to-end security
- Efficient resources usage
- Data processing
- …

Wireless Sensor Networks

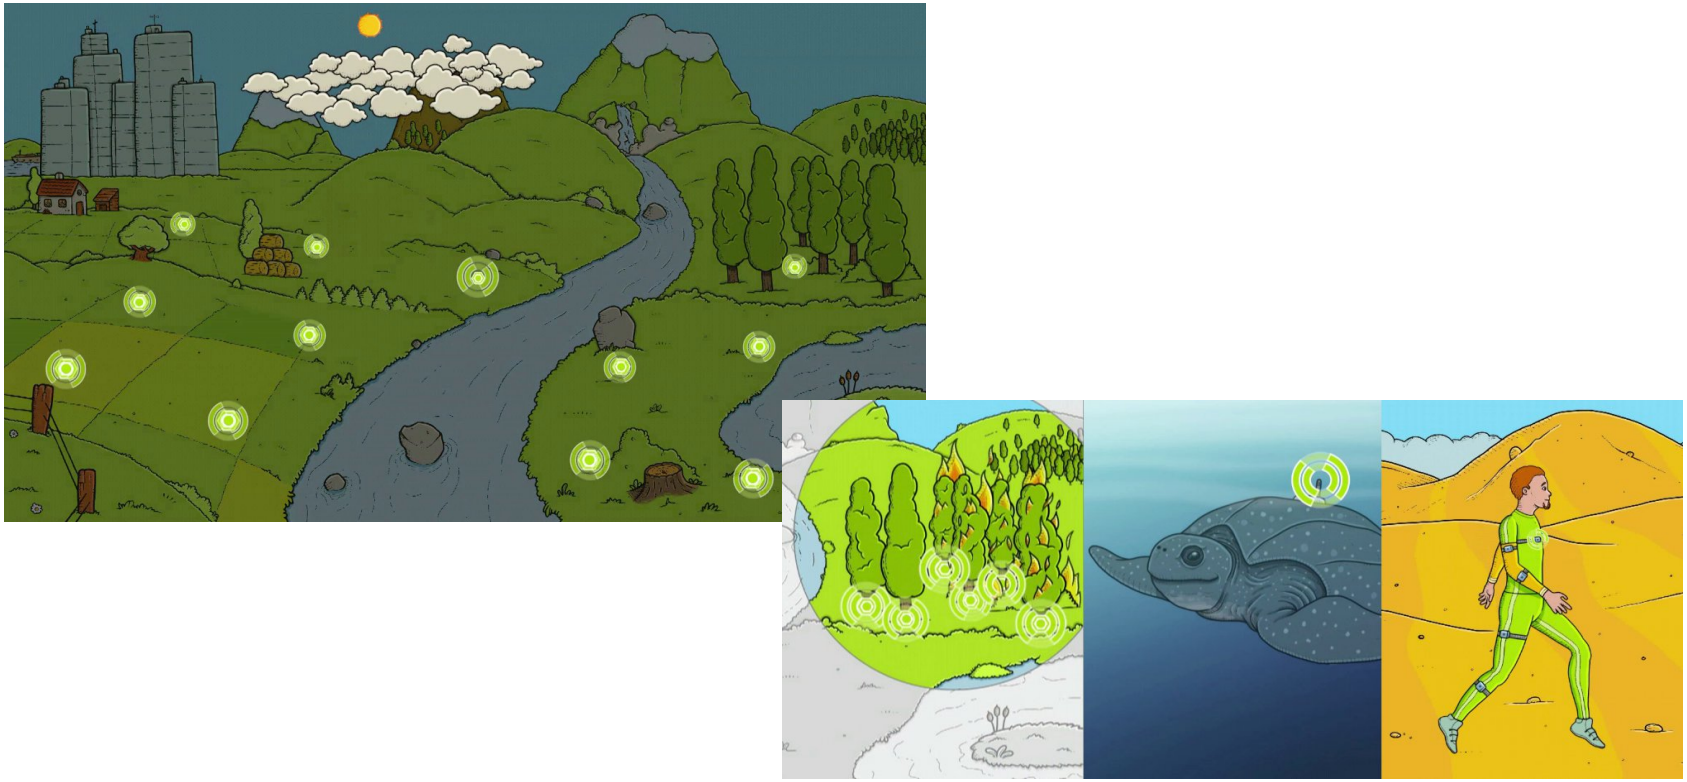

Internet of Things

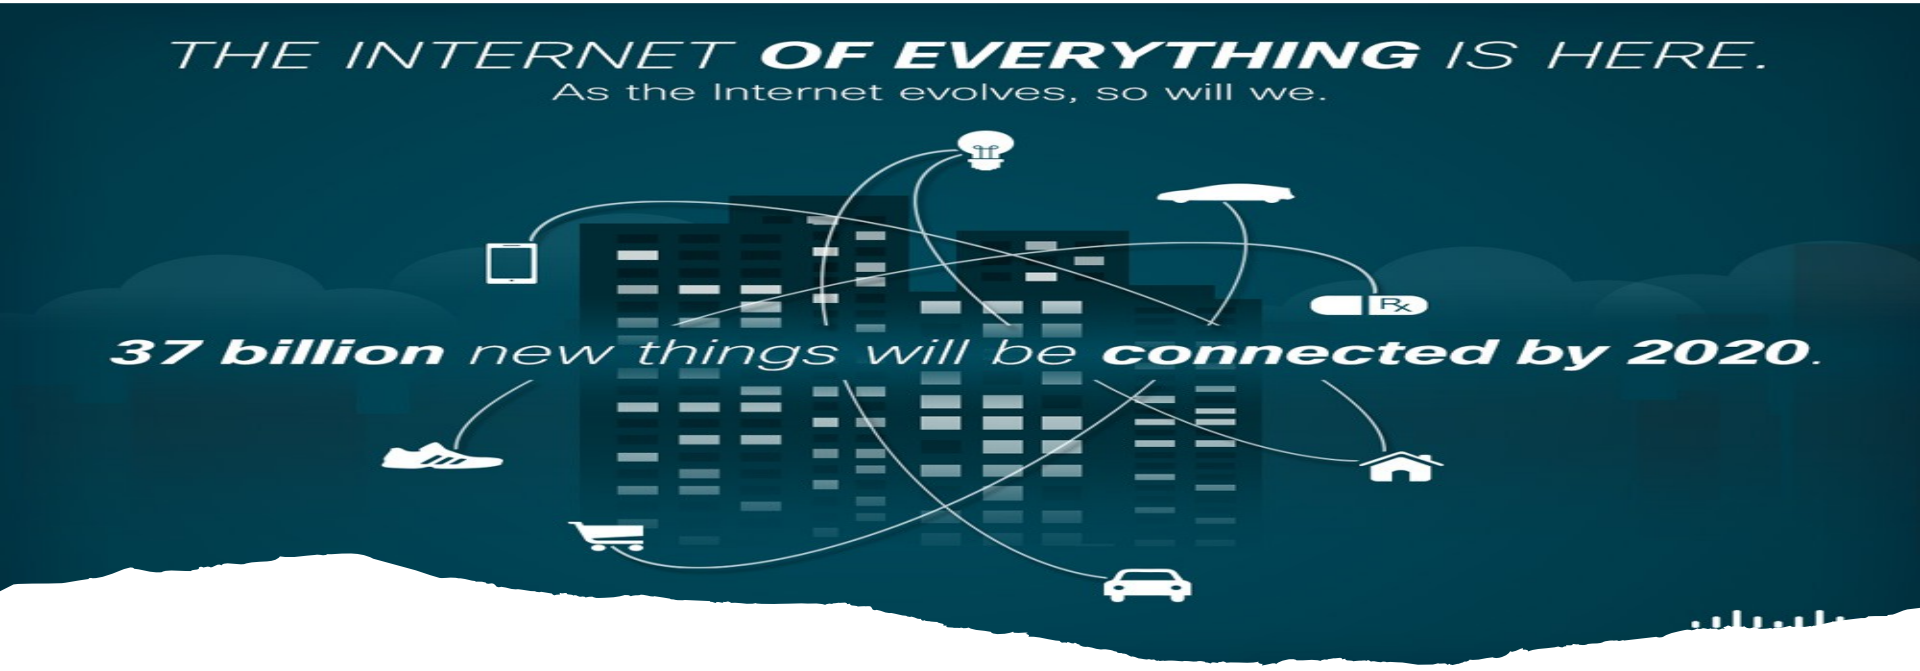

Datacenter & Cloud evolutions

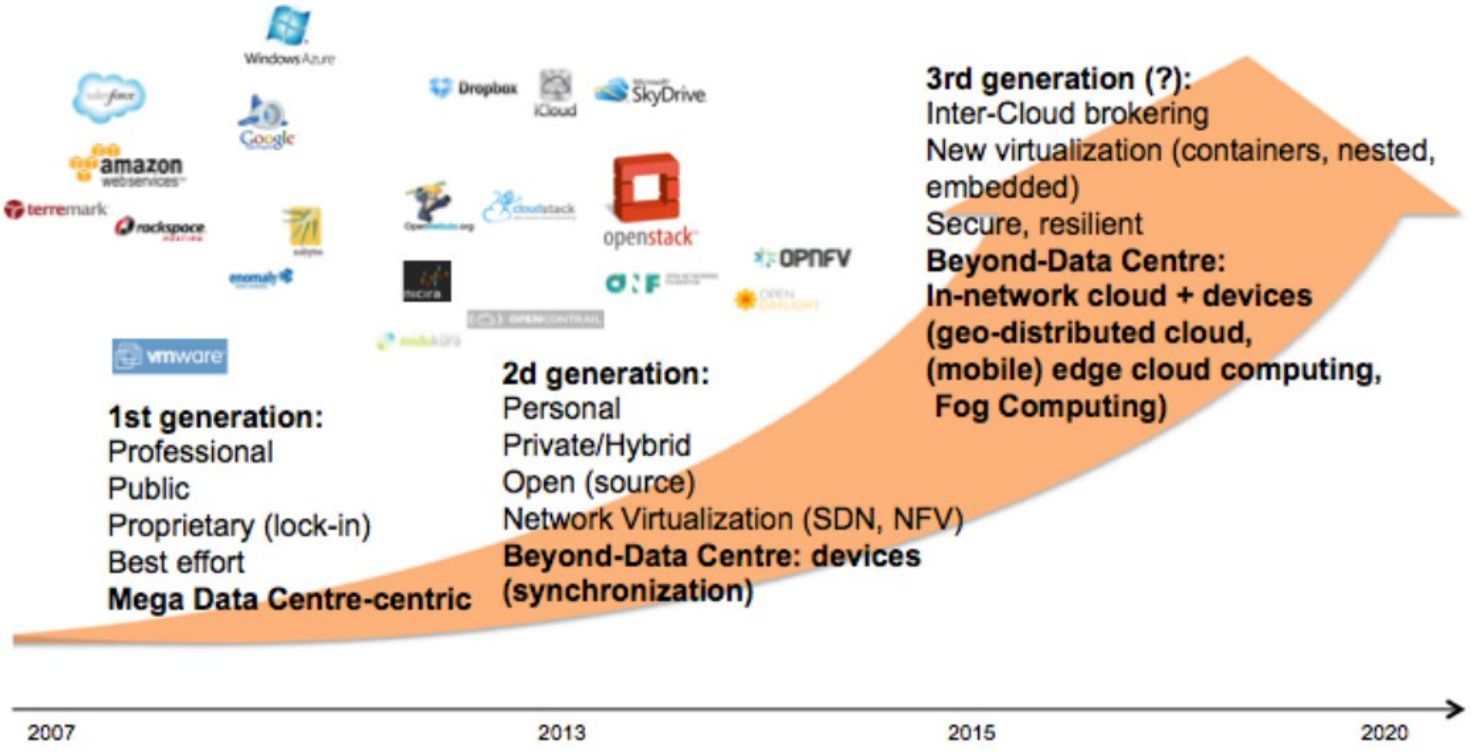

- (slides from 'Introduction to SILECS', F. Desprez, Apr. 19)

Convergence

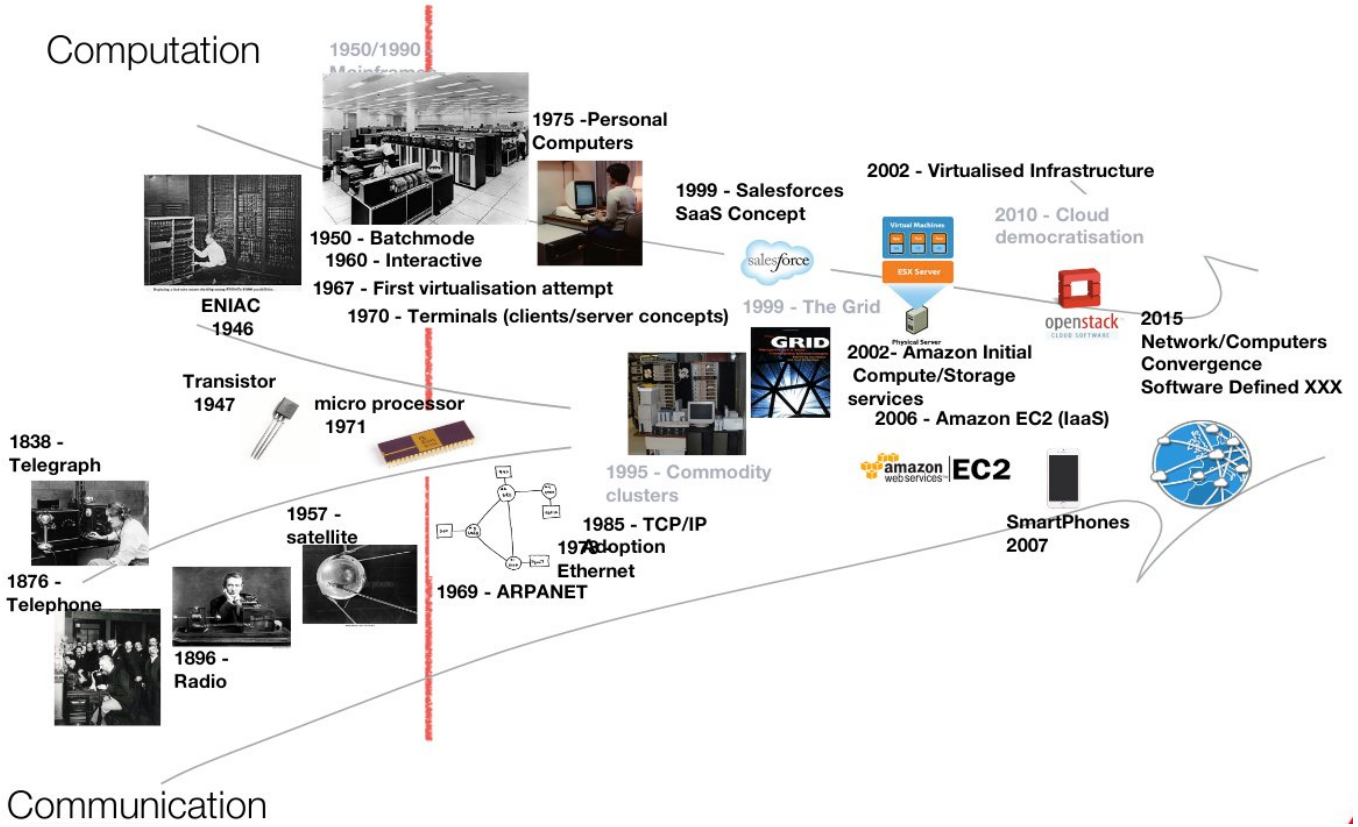

- (slides from ‘Introduction to SILECS’, F. Desprez, Apr. 19)

Research Cycle

- Idea
- Model
- Simulation
- **Experimentation**
- Deployment

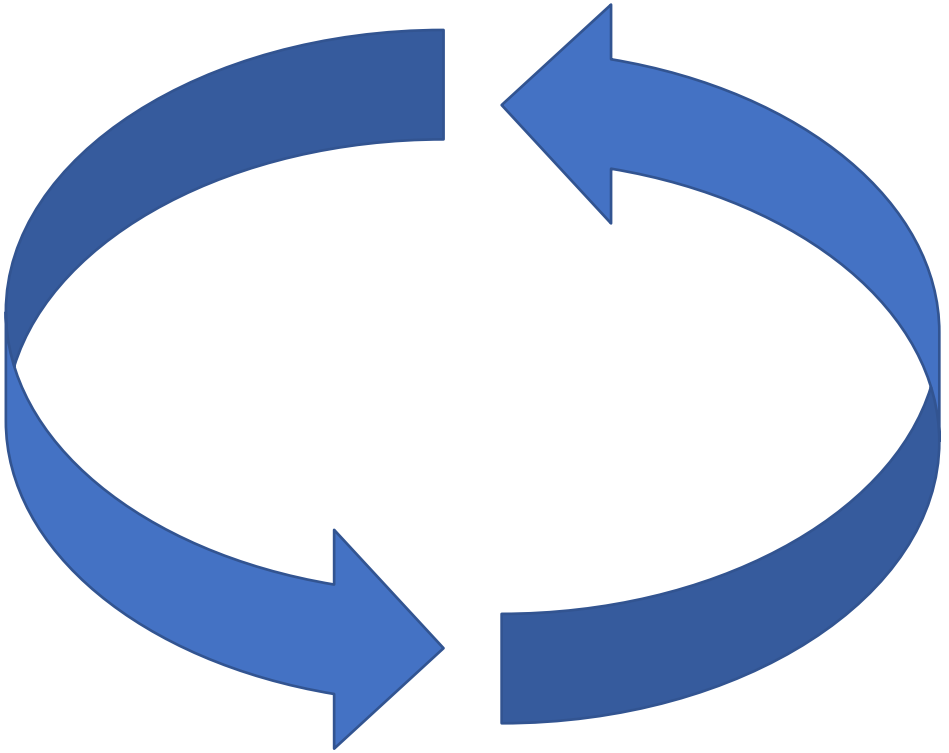

Association for Computing Machinery

How to experiment ?

- *SILECS*
  - *Super Infrastructure for Large-Scale Experimental Computer Science*
- Meta testbed for reproducible experimental research
- Merge of different testbeds
  - Grid'5000
  - FIT (IoT-LAB, CortexLab, …)
- French Label : Infrastructure de Recherche nationale
  - IR SILECS (2018)
- A common roadmap
  - Unified portail
  - Unified monitoring
  - Improve uses cases

From Sensors to the Cloud

- Complementary testbeds
  - IoT-LAB
  - Grid'5000
- Many use cases :
  - Edge/Fog Computing : Data generated in IoT-LAB and processed in G5K
  - Energy consumption of IoT deployments, from sensors to datacenters
  - LoRaWAN infrastructure: IoT-LAB for LoRaWAN nodes & gateway, G5K for LoRaWAN server
  - ...

IoT-LAB

SILECS Testbeds Overview

Experimentation

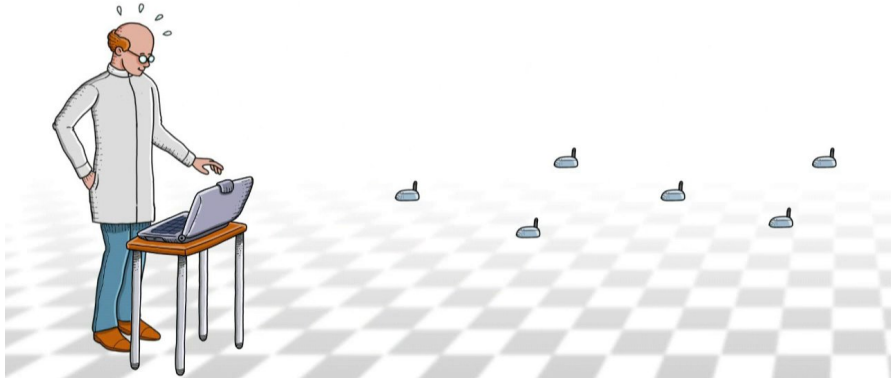

How to easily develop and test a **large scale** IoT application

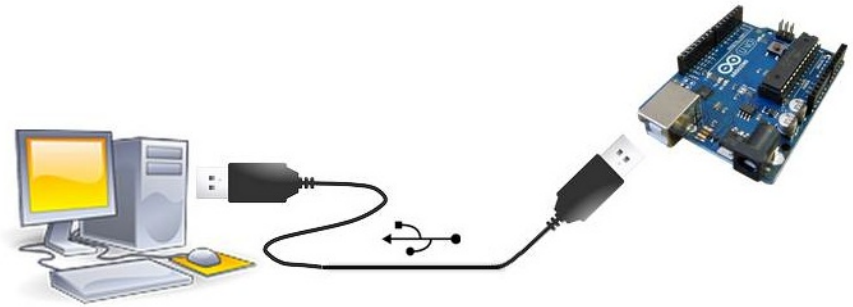

IoT-LAB

- The Very Large Scale IoT Testbed
  - 1500 IoT nodes
  - 7 sites located in France

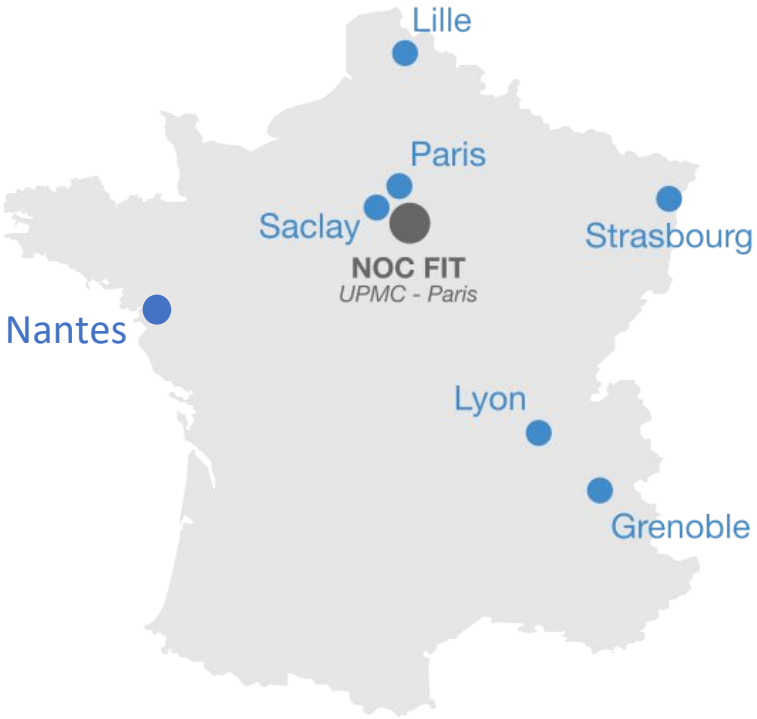

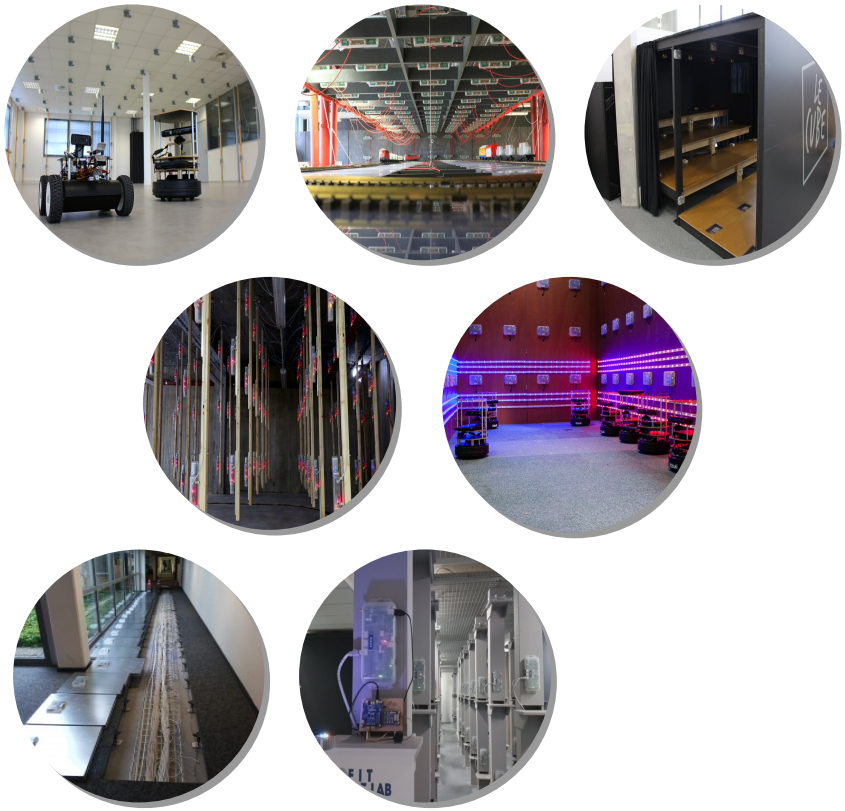

IoT-LAB

- Fully automated
  - Available 24/7
- Reproductible experimentations
- Multi-sites
- Free Access for everyone
  - Academic (researchers, students)
  - Industrials

Strasbourg site


FIT IoT LAB
LILLE - EURATECHNOLOGIES
LILLE - HAUTE BORNE
STRASBOURG
PARIS
ROCQUENCOURT
RENNES
7 SITES
2000 CAPTEURS SANS FIL

Lille site

IoT-LAB Nodes

- 3 logical components :
  - **Open Node** : programmable IoT Node
  - **Gateway** : uplink to the infrastructure, deploy user firmware
  - **Control Node** : monitoring (energy, radio)

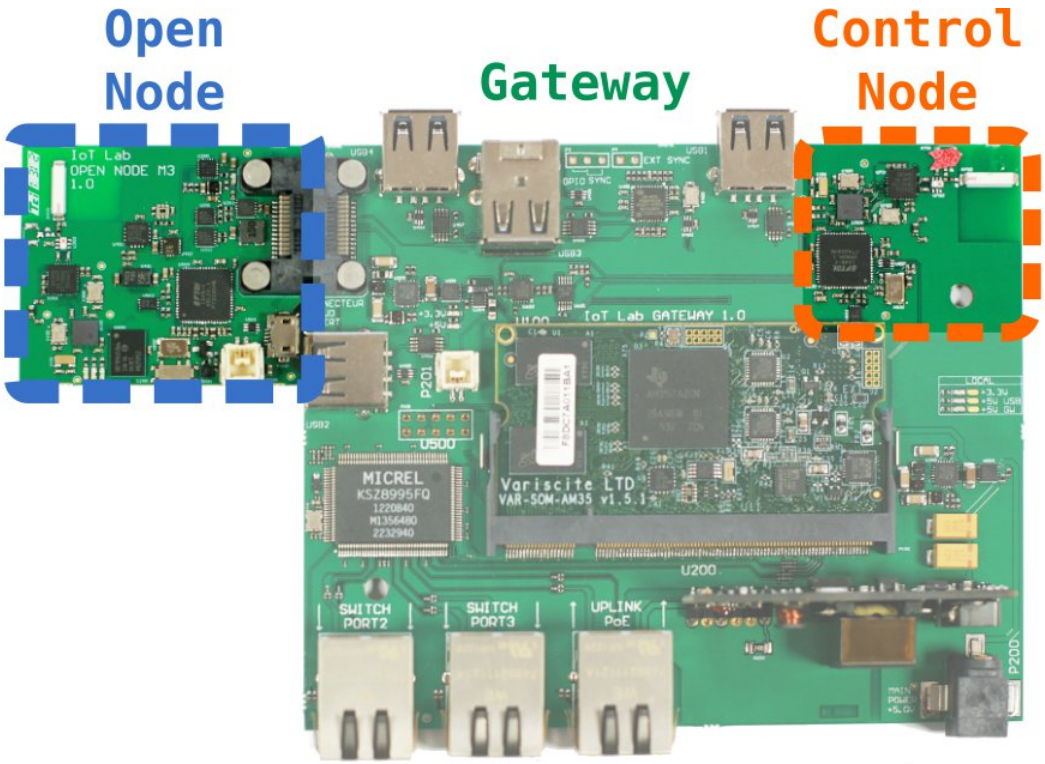

IoT-LAB Nodes (legacy)

Name (nb)	MCU	Sensors	Radio	
M3 (817)	Cortex M3 (32bits), 72 MHz, 256 kB ROM, 64 kB RAM	• Light • Accelerometer • Pression	• AT86RF231 (2.4GHz)	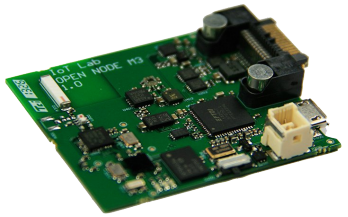
A8 (470)	Cortex A8 (32 bits), 600 Mhz, 256 MB RAM	• Light • Accelerometer • Pression	• AT86RF231 (2.4GHz) • Ethernet	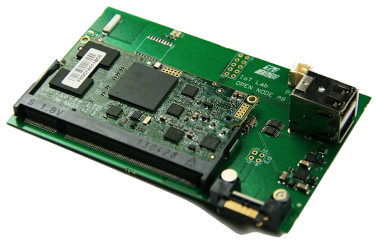

IoT-LAB Nodes (custom)

- Open Nodes from the market
  - Arduino-zero, Zolertia, ST, nRF, micro:bit, Pycom, etc.
  - https://www.iot-lab.info/docs/boards/overview/
- Requirements : USB interface + Linux toolchain

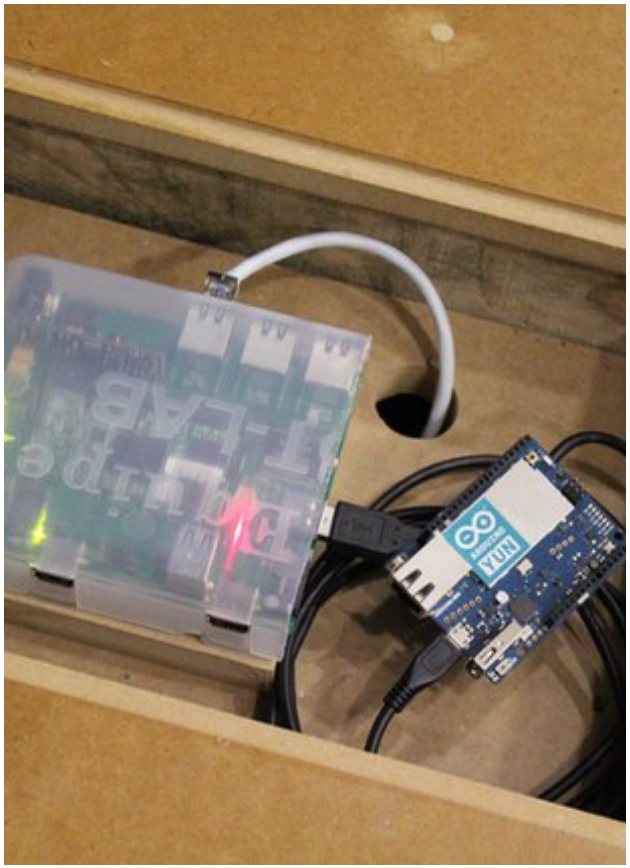

Embedded OS

OS	M3	A8	CUSTOM
FreeRTOS	X	-	X
Contiki	X	-	X
Riot	X	-	X
OpenWSN	X	-	X
Zephyr	-	-	X
Linux Yocto	-	X	-

- Test your own OS on our nodes !

Features

- Large scale user firmware deployment
- Automatic performances monitoring
  - Energy, radio level, radio capture, RTL-SDR
- Serial port and debug port access
- User workspace for development
  - Via remote server trough SSH
  - Via local virtual machine
- Public IPv6 networks
- LoRaWAN Infrastructure

Tools

- **Web Portal**: quick hands-on for beginners
- **CLI-tools + Run Script**: batch your experiment
  - experiment, node, profile
- **Serial_aggregator**: gather nodes serial output
- **OML Plot Tools**: graph monitoring results
- **Remote debugger**: gdb Open Node
- **Sniffer_aggregator**: gather radio capture
- **Radio characterization**: understand radio topology
- **MQTT & Leshan broker**: forward data to Internet

Global Overview

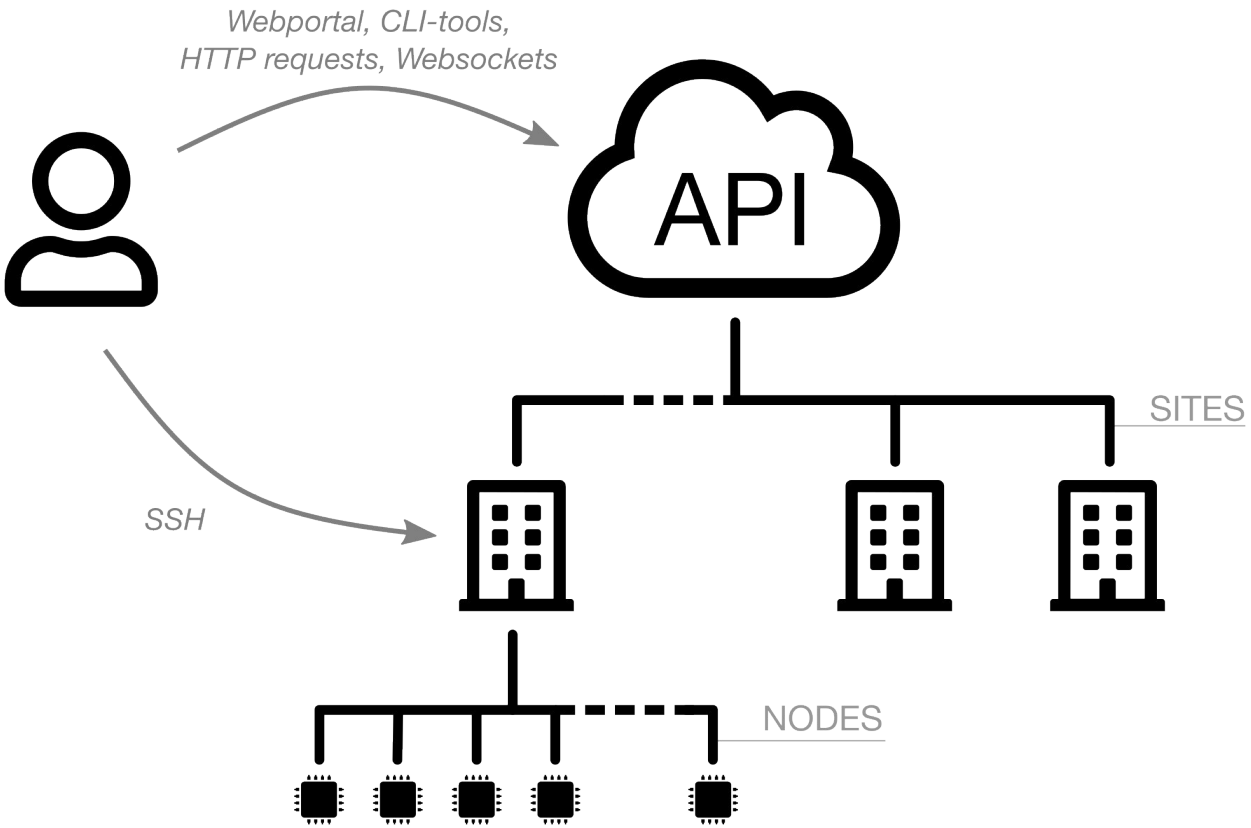

REST API

- Authenticated access
- Experimentation
  - Submit, reload, stop or cancel, resources descriptions, etc.
- Monitoring profile
  - Get, create, modify, delete
- User preferences
  - Modify user, SSH keys, password, etc.

Learn

- Quickly hands-on IoT-LAB, Jupyter Notebook & Tutorials
  - https://www.iot-lab.info/learn/

- MOOC : IoT with MCU: a hands-on course
  - https://www.fun-mooc.fr/courses/course-v1:inria+41020+session01/about

Community

- Official Site http://www.iot-lab.info
- Github https://github.com/iot-lab/
- Mailing list users@iot-lab.info

- Register https://www.iot-lab.info/testbed/signup

Grid’5000

SILECS Testbeds Overview

Grid’5000 overview

- A scientific instrument, targeting distributed computing research community (in a broad sense)

  ⇒ Cloud, HPC, Network, Datacenters, AI…
- Design goals:
  - A testbed for experimental research (even at a large scale)
  - Provide reconfigurable and observable environment
  - Targets high-quality and reproducible environment
- Exists for 15 years: 600+ active users each year, 100-150 scientific publications, 50M+ cores.hours used

Contributors

- Supported by main French research institutions
- Scientific committees that drive evolutions
- Technical team that runs the platform (~10 people)

Infrastructure

- **A large-scale infrastructure**
  - 8 sites, 39 clusters, ±800 nodes,
  - ±16000 CPU cores and ±300 GPU
  - Dedicated 10-Gbps backbone network
- **Wide variety of hardware**
  - Several generation of Intel & AMD CPUs
  - ARM64 and Power8 also available
  - GPUs: Nvidia (from Geforce to A100) but also AMD
  - Infiniband and Omnipath networks, High speed Ethernet, Persistent Memory, SSD/NVMe…

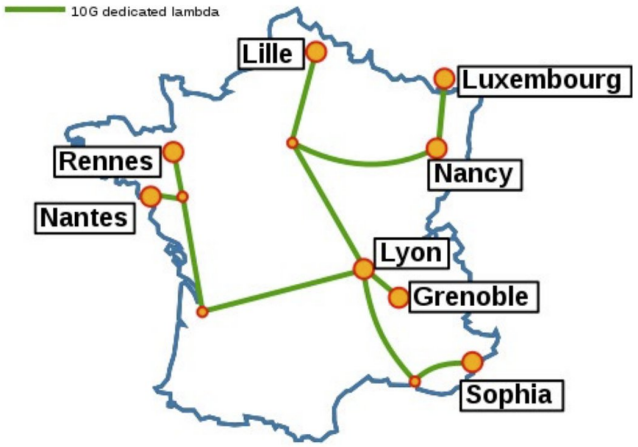

Discovering resources

https://api.grid5000.fr/stable/sites/lyon/clusters/nova/nodes/nova-1?pretty


```
{
  "architecture": {
    "nb_cores": 16,
    "nb_procs": 2,
    "nb_threads": 32,
    "platform_type": "x86_64"
  },
  "bios": {
    "configuration": {
      "cstate_c1e": true,
      "cstate_enabled": true,
      "ht_enabled": true,
      "turboboost_enabled": true
    },
    "release_date": "09/08/2016",
    "vendor": "Dell_Inc.",
    "version": "2.2.5"
  },
  "chassis": {
    "manufacturer": "Dell_Inc.",
    "name": "PowerEdge_R430",
(...)
```


- **nodes**, grouped in **clusters**, distributed over Grid'5000 **sites**
  - *nova-5.lyon.grid5000.fr*
- Check Hardware page on website
- Or **reference API** for complete nodes description

Reserving resources

- OAR Resource Management System
  - used to submit “job” experiment / reserving resource in advance

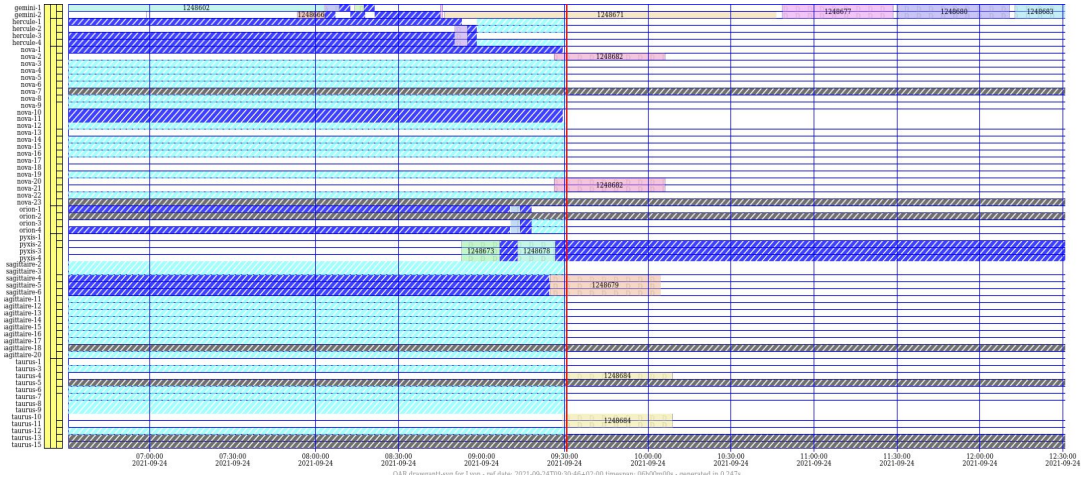

- Usage policy :
  - No more than 2 hours of a full cluster during the day (or half of the cluster for 4 hours, etc.). No restriction during night and weekend. Exemption possible if motivated.
  - Promotes interactive usage (experiment design) during days, experiments runs at night

Experiment tools

- Software Installation & Low level system configuration
  - *sudo-g5k* to become root on the default OS environment
  - *Kadeploy* to deploy your own OS
- Network related:
  - *Kavlan* to isolate nodes in your own LAN
  - Emulated network over Grid'5000
  - IPv6
- Reservation of individual disks inside nodes (and several other storage solutions)
- Tools to spawn VM/containers
- HPC-friendly environment
- …

Running experiment, monitoring

- It's up to the user to design its experiments
  - Usually, first trials needs connecting to resources interactively using SSH
  - Then, use scripts (from bash scripts to high level programming)

    ⇒ more on that later
- **Monitoring** and getting results :

  No single answer

  We provide *Kwollect* monitoring framework :
  - focused on "environmental" metrics (i.e. not available from inside the nodes : power consumption, temperature, network devices, etc.)
  - also prometheus-based metrics on nodes (even custom ones defined by users)
  - Metrics values are available using an API and can be visualized on a dashboard

Running experiment, monitoring


```
sdelamare@taurus-3(15:19):~$ curl "https://api.grid5000.fr/stable/sites/lyon/metrics?job_id=1248323"
[{"timestamp":"2021-09-23T15:17:45+02:00","device_id":"taurus-3","metric_id":"wattmetre_power_watt","value":141.514,"labels":{"_device_orig": ["wattmetre4-port2"]}},
 {"timestamp":"2021-09-23T15:17:45.14901+02:00","device_id":"taurus-3","metric_id":"bmc_ambient_temp_celsius","value":18,"labels":{}},
 {"timestamp":"2021-09-23T15:17:45.14901+02:00","device_id":"taurus-3","metric_id":"bmc_node_power_watt","value":98,"labels":{}},
 {"timestamp":"2021-09-23T15:17:46+02:00","device_id":"taurus-3","metric_id":"wattmetre_power_watt","value":125.34,"labels":{"_device_orig": ["wattmetre4-port2"]}},
 {"timestamp":"2021-09-23T15:17:47+02:00","device_id":"taurus-3","metric_id":"wattmetre_power_watt","value":120.583333333333,"labels":{"_device_orig": ["wattmetre4-port2"]}},
 {"timestamp":"2021-09-23T15:17:47.127532+02:00","device_id":"taurus-3","metric_id":"network_ifacein_bytes_total","value":2963739631,"labels":{"interface": "eth0", "_device_orig": ["force10-port-3"]}},
 {"timestamp":"2021-09-23T15:17:47.127532+02:00","device_id":"taurus-3","metric_id":"network_ifaceout_packets_total","value":2109310677,"labels":{"interface": "eth0", "_device_orig": ["force10-port-3"]}},
 {"timestamp":"2021-09-23T15:17:47.127532+02:00","device_id":"taurus-3","metric_id":"network_ifaceout_bytes_total","value":1575825569,"labels":{"interface": "eth0", "_device_orig": ["force10-port-3"]}},
 {"timestamp":"2021-09-23T15:17:47.127532+02:00","device_id":"taurus-3","metric_id":"network_ifacein_packets_error_total","value":0,"labels":{"interface": "eth0", "_device_orig": ["force10-port-3"]}},
 {"timestamp":"2021-09-23T15:17:47.127532+02:00","device_id":"taurus-3","metric_id":"network_ifacein_packets_total","value":0,"labels":{"interface": "eth0", "_device_orig": ["force10-port-3"]}},
 {"timestamp":"2021-09-23T15:17:47.127532+02:00","device_id":"taurus-3","metric_id":"network_ifacein_packets_discard_total","value":0,"labels":{"interface": "eth0", "_device_orig": ["force10-port-3"]}},
 {"timestamp":"2021-09-23T15:17:47.127532+02:00","device_id":"taurus-3","metric_id":"network_ifaceout_packets_error_total","value":0,"labels":{"interface": "eth0", "_device_orig": ["force10-port-3"]}},
 {"timestamp":"2021-09-23T15:17:47.127532+02:00","device_id":"taurus-3","metric_id":"network_ifaceout_packets_discard_total","value":0,"labels":{"interface": "eth0", "_device_orig": ["force10-port-3"]}},
 {"timestamp":"2021-09-23T15:17:48+02:00","device_id":"taurus-3","metric_id":"wattmetre_power_watt","value":100.229166666667,"labels":{"_device_orig": ["wattmetre4-port2"]}},
 {"timestamp":"2021-09-23T15:17:49+02:00","device_id":"taurus-3","metric_id":"wattmetre_power_watt","value":98.5734693877551,"labels":{"_device_orig": ["wattmetre4-port2"]}},
 {"timestamp":"2021-09-23T15:17:50+02:00","device_id":"taurus-3","metric_id":"wattmetre_power_watt","value":114.370833333333,"labels":{"_device_orig": ["wattmetre4-port2"]}},
 {"timestamp":"2021-09-23T15:17:50.148934+02:00","device_id":"taurus-3","metric_id":"bmc_node_power_watt","value":98,"labels":{}},
 {"timestamp":"2021-09-23T15:17:50.148934+02:00","device_id":"taurus-3","metric_id":"bmc_ambient_temp_celsius","value":18,"labels":{}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0.54,"labels":{"cpu": "18", "mode": "user"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":2815.83,"labels":{"cpu": "19", "mode": "idle"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":1.01,"labels":{"cpu": "19", "mode": "iowait"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "19", "mode": "irq"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "19", "mode": "nice"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "19", "mode": "softirq"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "19", "mode": "steal"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":1.3,"labels":{"cpu": "19", "mode": "system"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":1.96,"labels":{"cpu": "19", "mode": "user"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":2816.74,"labels":{"cpu": "2", "mode": "idle"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0.72,"labels":{"cpu": "2", "mode": "iowait"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "13", "mode": "irq"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "2", "mode": "irq"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "2", "mode": "nice"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0.09,"labels":{"cpu": "2", "mode": "softirq"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "2", "mode": "steal"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":1.48,"labels":{"cpu": "2", "mode": "system"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":1.01,"labels":{"cpu": "2", "mode": "user"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":2812.25,"labels":{"cpu": "20", "mode": "idle"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":6.2,"labels":{"cpu": "20", "mode": "iowait"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "20", "mode": "irq"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "20", "mode": "nice"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0.01,"labels":{"cpu": "20", "mode": "softirq"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "20", "mode": "steal"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":1.24,"labels":{"cpu": "20", "mode": "system"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0.62,"labels":{"cpu": "20", "mode": "user"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":2817.67,"labels":{"cpu": "21", "mode": "idle"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0.7,"labels":{"cpu": "21", "mode": "iowait"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "21", "mode": "irq"}},
 {"timestamp":"2021-09-23T15:17:50.283066+02:00","device_id":"taurus-3","metric_id":"prom_node_cpu_seconds_total","value":0,"labels":{"cpu": "21", "mode": "nice"}},
```

Running experiment, monitoring

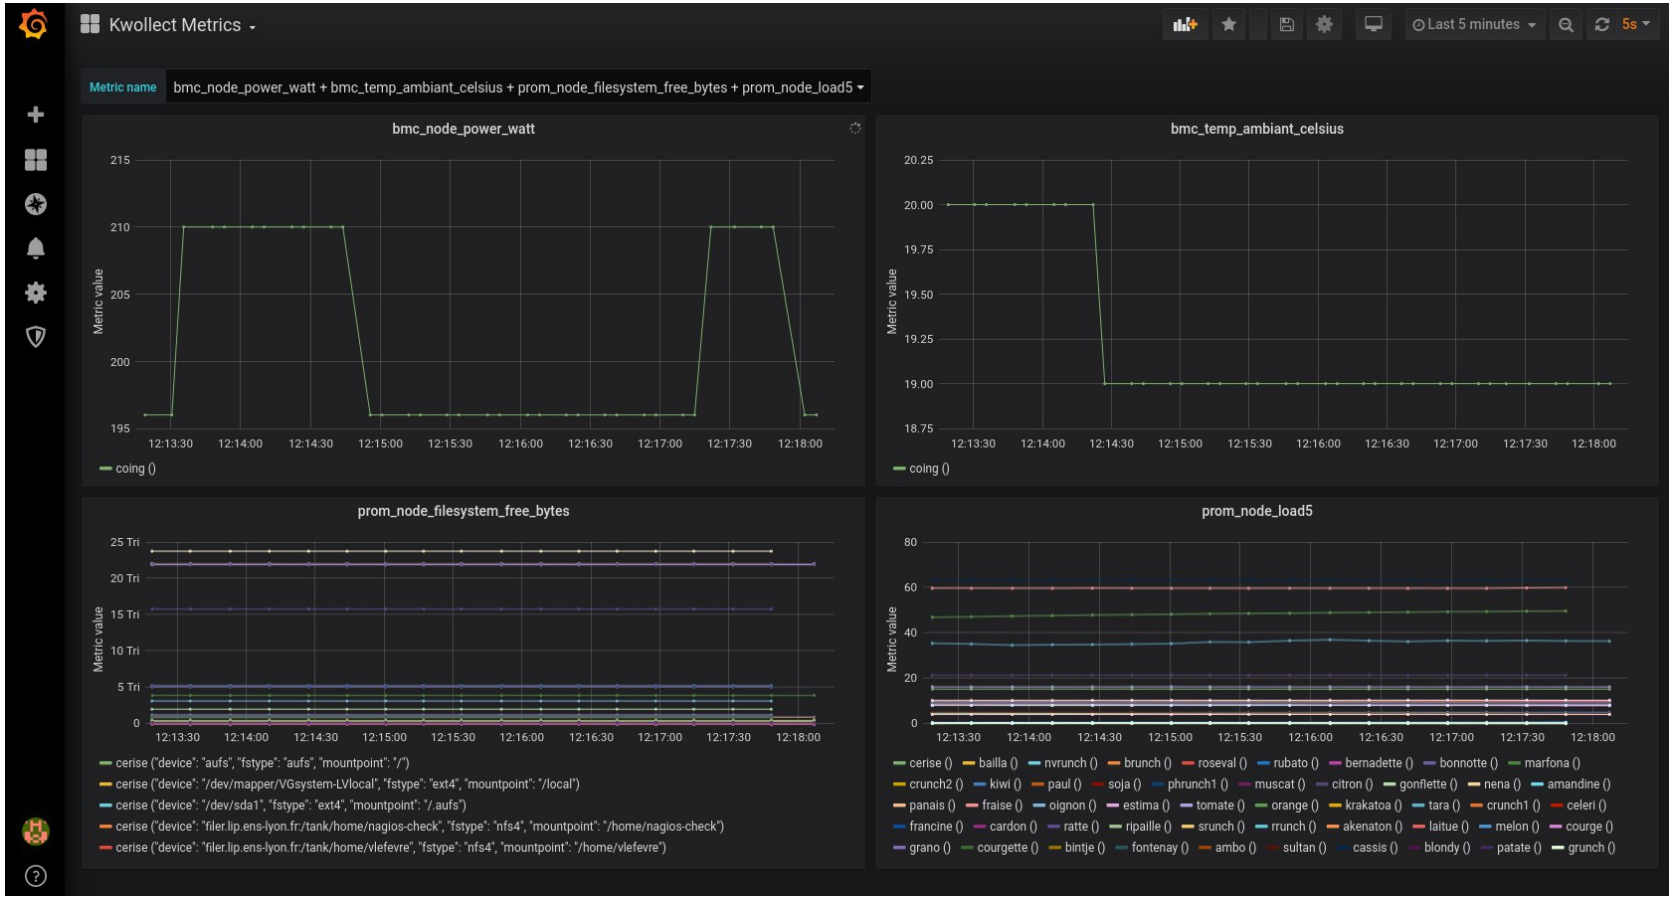

Experiment automation

- Every Grid'5000 tools can be manipulated using *Grid'5000 REST API*

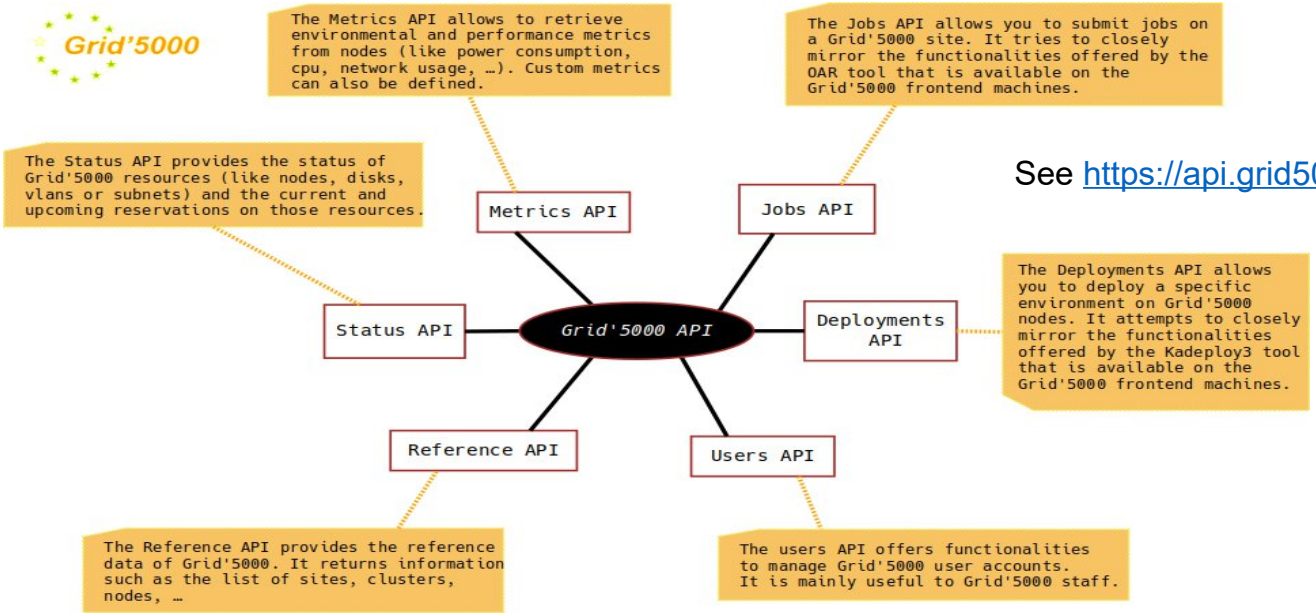

See https://api.grid5000.fr/doc/3.0/

- Base layer of several libraries and tools for experiment scripting and automation (contributed by Grid'5000 users)
  - *enoslib, execo, python-grid5000, ruby-cute…*
  - Meet different needs: low vs high level, only focused on Grid'5000 or generalist tool for experiments, etc.

About Reproducible Experiments

Reproducible environment?

- IoT-Lab & Grid'5000 testbeds intended for experiment reproducibility
  - Re-use same nodes/firmware/OS/software across experiments
- Both testbeds offer a trusted environment:
  - We want results obtained by users to be reliable
  - Correct functioning of the infrastructure is carefully checked:
    - nodes continuously verified for HW error, CI to catch regressions on services offered to users
  - Evolutions in infrastructure are documented:
    - changes on hardware (Reference API's) and software (Environment Recipes) are tracked on Git repositories

Reproducible environment?

- But some outside perturbation cannot be controlled!
  - Electromagnetic environment in IoT-LAB (e.g. Wi-Fi usage in the campus)
  - Even in Grid'5000: datacenter temperature may cause CPU throttling, network utilization by others experiments on shared switches / backbone network, etc.

Packet loss
During workhours

- Possible solutions:
  - Schedule experiments when perturbation are less likely (e.g. outside office hours for Wi-Fi usage)
  - Reserve more resources than really needed (e.g. all nodes connected to a switch)
  - Safeguards on testbeds sides : anechoic chamber, more dedicated/reservable resources

(Reproducible) experiment outline

1/ Discovering and selecting resources

2/ Reconfiguring resources according to experiment needs

3/ Designing the experiment, the monitoring and collection of results

4/ Automating the experiment workflow

- Both testbeds provide API for experimentation scripting

  ⇒ essential to reproducibility

  - w/ support in enoslib

Conclusion

Conclusion (1/2)

- SILECS: still distinct testbeds today
- How to ease life of users who need resources from both testbeds ?
- We focused on network connectivity so far :
  - IPv6 on G5K w/ configurable firewall for direct connections with IoT-LAB resources
  - Under consideration: dedicated links between IoT-LAB & G5K existing in the same city.
- In the future, fill the gap gradually:
  - Many ideas: users account, monitoring framework, common APIs…

Conclusion (2/2)

- IoT-LAB & Grid'5000 : allow end to end experiment
- Designed with reproducibility in mind
- A lot to learn during this school!

Questions ?