SILECS

Super Infrastructure for Large-scale Experimental Computer Science

Serge Fdida (UPMC)

Frédéric Desprez (Inria)

Christian Perez (Inria)

INRIA, CNRS, RENATER, CEA, CPU, CDEFI, IMT, Sorbonne Université, Université Strasbourg, Université Lorraine, Université Grenoble Alpes, Université Lille 1, Université Rennes 1, Université Toulouse, ENS Lyon, INSA Lyon

Introduction

- Exponential improvement of
  - Electronics (energy consumption, size, cost)
  - Capacity of networks (WAN, wireless, new technologies)
- Exponential growth of applications near users
  - Smartphones, tablets, connected devices, sensors, …
  - Prediction of 50 billions of connected devices by 2020 (CISCO)

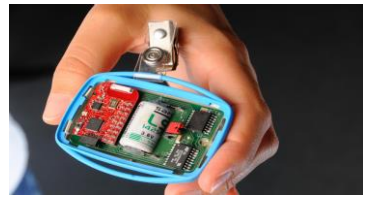

- Large number of Cloud facilities to cope with generated data
  - Many platforms and infrastructures available around the world
  - Several offers for IaaS, PaaS, and SaaS platforms
  - Public, private, community, and hybrid clouds
  - Going toward distributed Clouds (FOG, Edge, extreme Edge)

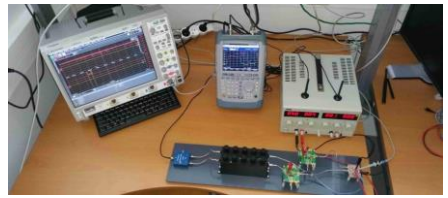

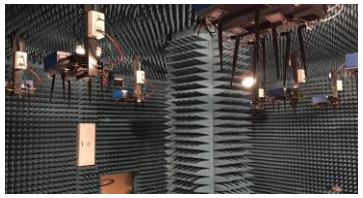

Convergence

Computation

1950/1990 - Mainframes

1975 -Personal Computers

ENIAC
1946

1950 - Batchmode
1960 - Interactive

1967 - First virtualisation attempt

1970 - Terminals (clients/server concepts)

1999 - Salesforces SaaS Concept

2002 - Virtualised Infrastructure

2010 - Cloud democratization

1999 - The Grid

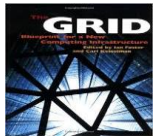

2002- Amazon Initial Compute/Storage services

2006 - Amazon EC2 (IaaS)

2015
Network/Computers Convergence
Software Defined XXX

Transistor
1947

micro processor
1971

1838 - Telegraph

1876 - Telephone

1957 - satellite

1896 - Radio

1969 - ARPANET

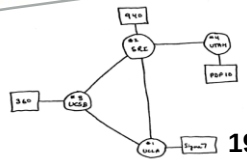

1985 - TCP/IP Adoption
1973 - Ethernet

1995 - Commodity clusters

SmartPhones
2007

Communication

Good experiments

A **good experiment** should fulfill the following properties

- **Reproducibility**: *must* give the same result with the same input
- **Extensibility:** *must* target possible comparisons with other works and extensions (more/other processors, larger data sets, different architectures)
- **Applicability:** *must* define realistic parameters and *must* allow for an easy calibration
- **"*Revisability*":** when an implementation does not perform as expected, *must* help to identify the reasons

SILECS: based upon two existing infrastructures

- **FIT**
  - Providing Internet players access to a variety of fixed and mobile technologies and services, thus accelerating the design of advanced technologies for the Future Internet
  - 4 key technologies and a single control point: IoT-Lab (connected objects & sensors, mobility), CorteXlab (Cognitive Radio), wireless (anechoic chamber), Network Operations Center, Advanced Cloud technology including OpenStack
  - 9 sites (Paris (2), Evry, Rocquencourt, Lille, Strasbourg, Lyon, Grenoble, Sophia Antipolis)
- **Grid'5000**
  - A scientific instrument for experimental research on large future infrastructures: Clouds, datacenters, HPC Exascale, Big Data infrastructures, networks, etc.
  - 10 sites, > 8000 cores, with a large variety of network connectivity and storage access, dedicated interconnection network granted and managed by RENATER
- **Software stacks dedicated to experimentation**
  - Resource reservation, disk image deployment, monitoring tools, data collection and storage

FIT

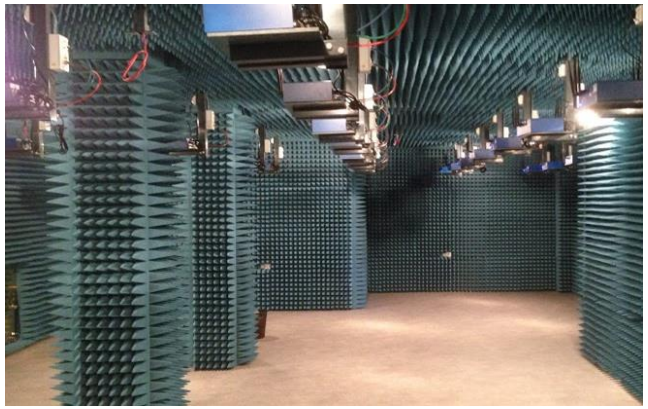

FIT-R2Lab: WiFi mesh testbed (DIANA)

FIT-CorteXlab: Cognitive Radio Testbed 40 Software Defined Radio Nodes (SOCRATE)

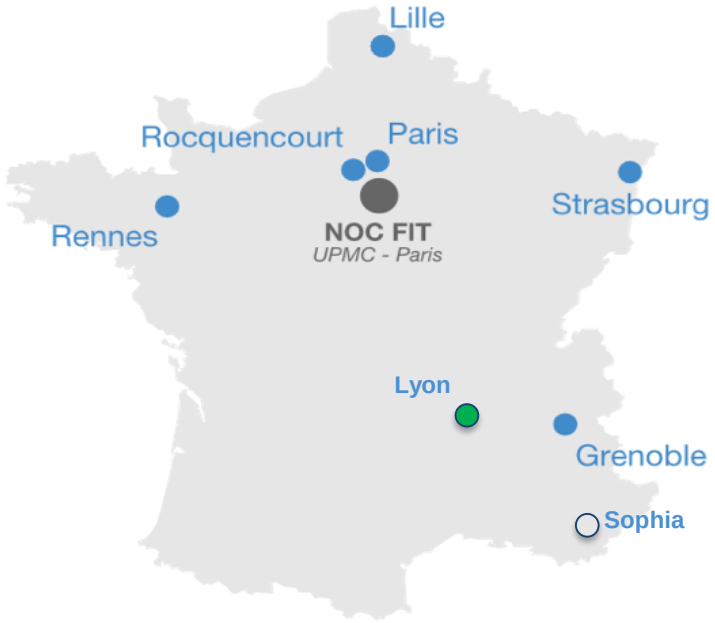

FIT-IoT-LAB

- 2700 wireless sensor nodes spread across six different sites in France
- Nodes are either fixed or mobile and can be allocated in various topologies throughout all sites

GRID'5000

- **Testbed for research on distributed systems**
  - Born from the observation that we need a better and larger testbed
  - HPC, Grids, P2P, and now Cloud computing and BigData systems
  - A complete access to the nodes' hardware in an exclusive mode (from one node to the whole infrastructure)
  - Dedicated network (RENATER)
  - Reconfigurable: nodes with Kadeploy and network with KaVLAN
- **Current status**
  - **10 sites, 29 clusters, 1060 nodes, 10474 cores, 10 Gbs network backbone**
  - Diverse technologies/resources (Intel, AMD, Myrinet, Infiniband, two GPU clusters, energy probes)
- **Some Experiments examples**
  - In Situ analytics
  - Big Data Management
  - HPC Programming approaches
  - Network modeling and simulation
  - Energy consumption evaluation
  - Batch scheduler optimization
  - Large virtual machines deployments

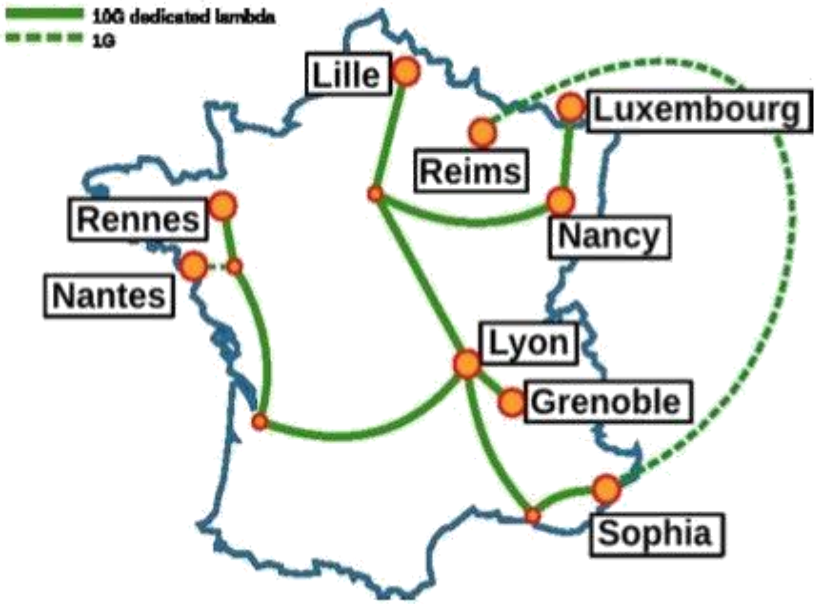

https://www.grid5000.fr/

Silecs: Envisioned Architecture

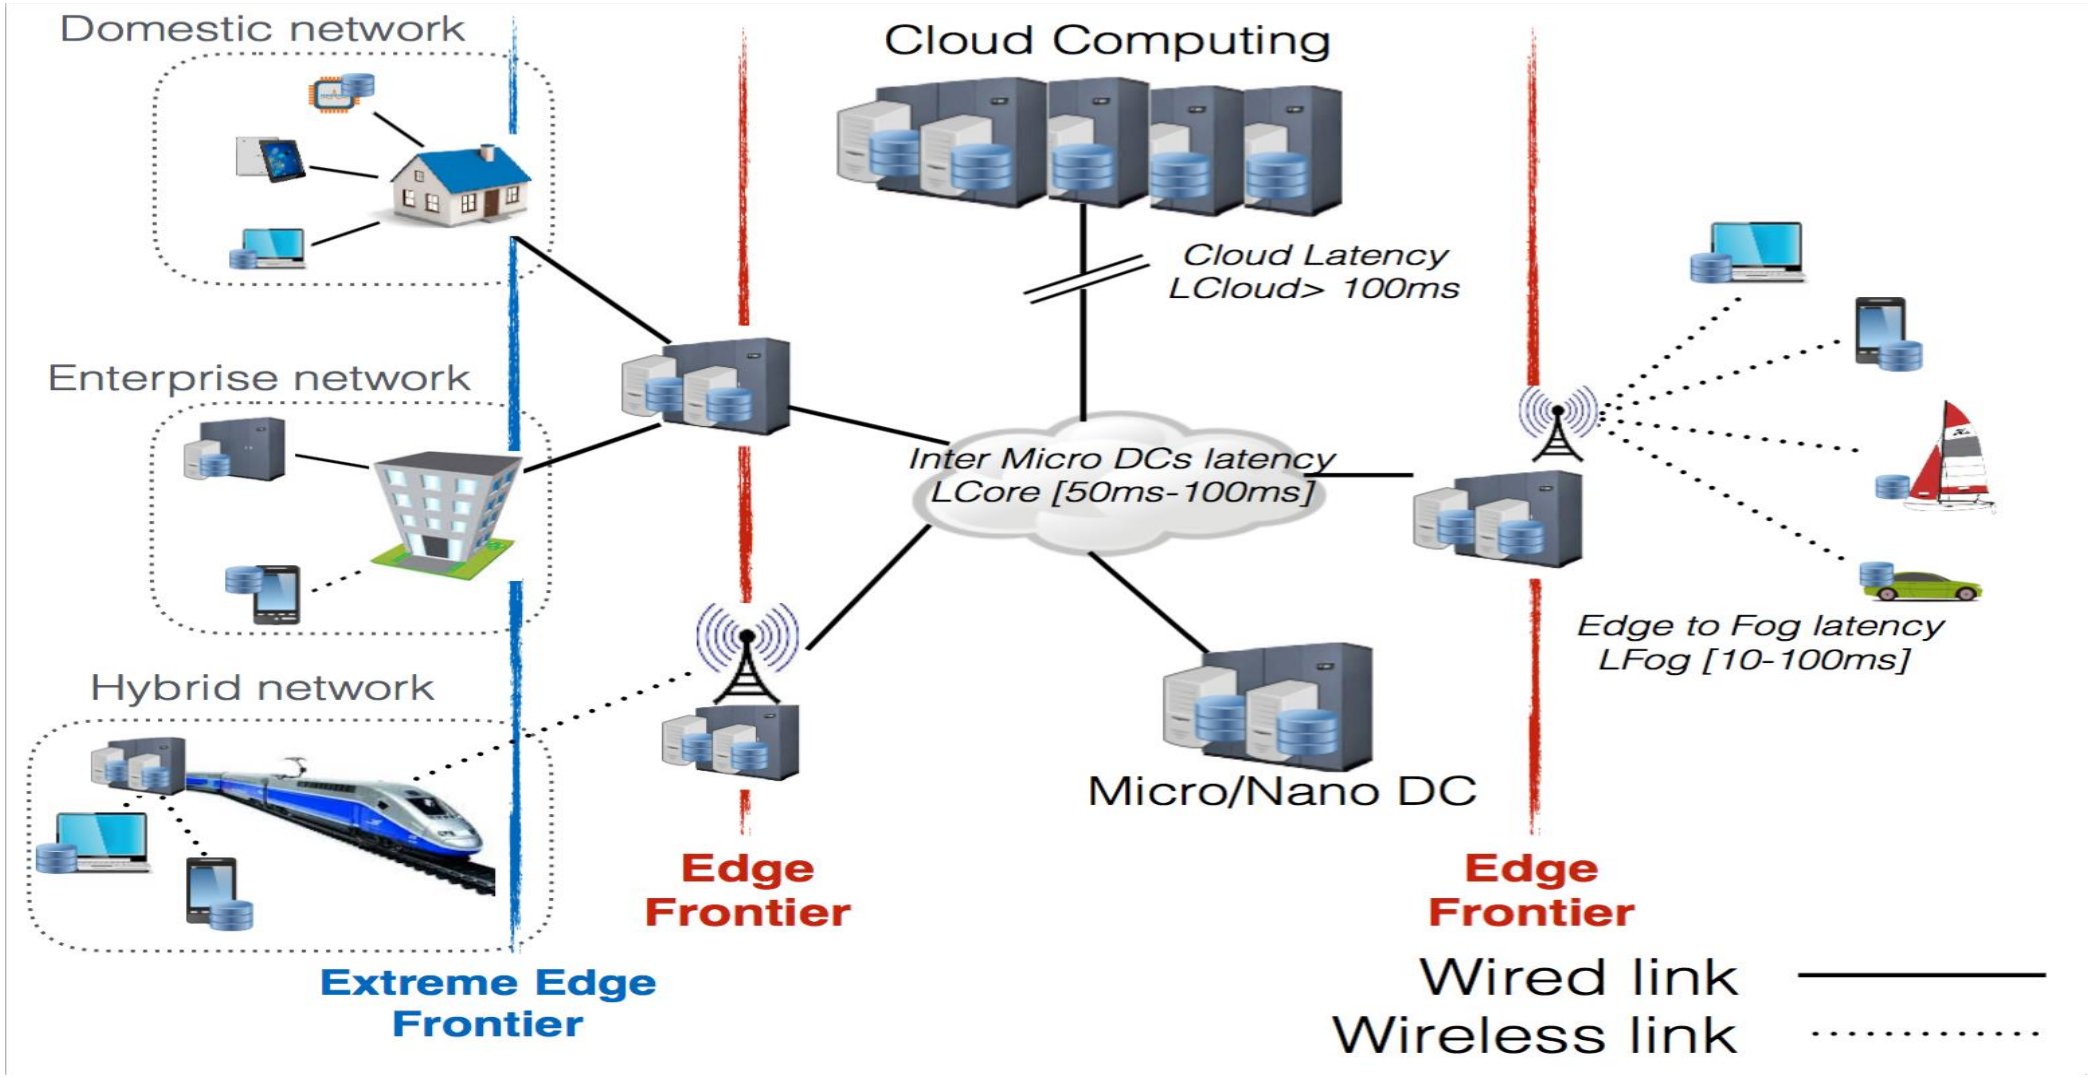

Credits: Bastien Confais

Silecs: Short Term View of the Architecture

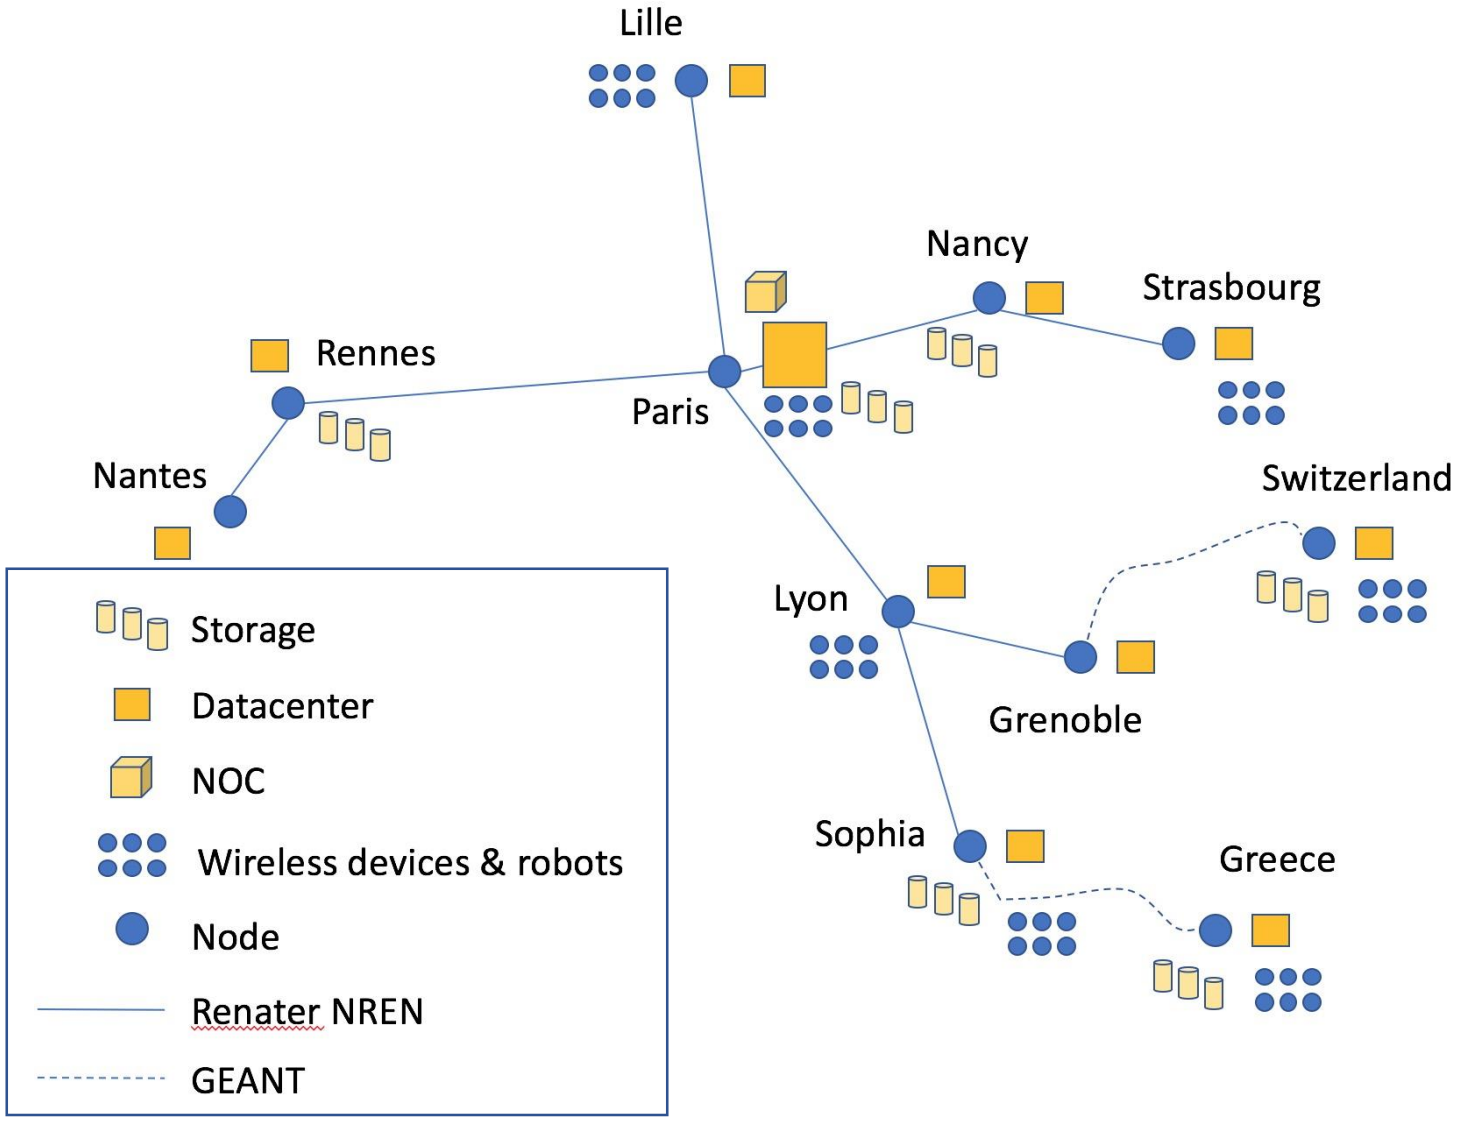

Data Center Portfolio

Targets

- Performance, resilience, energy-efficiency, security in the context of data-center design, Big Data processing, Exascale computing, etc.

Hardware

- Servers: x86, ARM64, POWER, accelerators (GPU, FPGA)
- Networking: Ethernet (10G, 40G), HPC networks (InfiniBand, Omni-Path)
- Storage: HDD, SSD, NVMe, both in storage arrays and clusters of servers

Experimental support

- Bare-metal reconfiguration
- Large clusters
- Integrated monitoring (performance, energy, temperature, network traffic)

Wireless Portfolio

Targets

- Performance, security, safety and privacy-preservation in complex sensing environment,
- Performance understanding and enhancement in wireless networking,
- Target applications: smart cities/manufacturing, building automation, standard and interoperability, security, energy harvesting, health care.

Hardware

- Software Defined Radio (SDR), LTE-Advanced and 5G
- Wireless Sensor Network (WSN/IEEE 802.15.4), LoRa/LoRaWAN
- Wifi/WIMAX (IEEE 802.11/16)

Experimental support

- Bare-metal reconfiguration
- Large-scale deployment (both in terms of densities and network diameter)
- Different topologies with indoor/outdoor locations
- Mobility-enabled with customized trajectories
- Anechoic chamber
- Integrated monitoring (power consumption, radio signal, network traffic)

The GRAIL

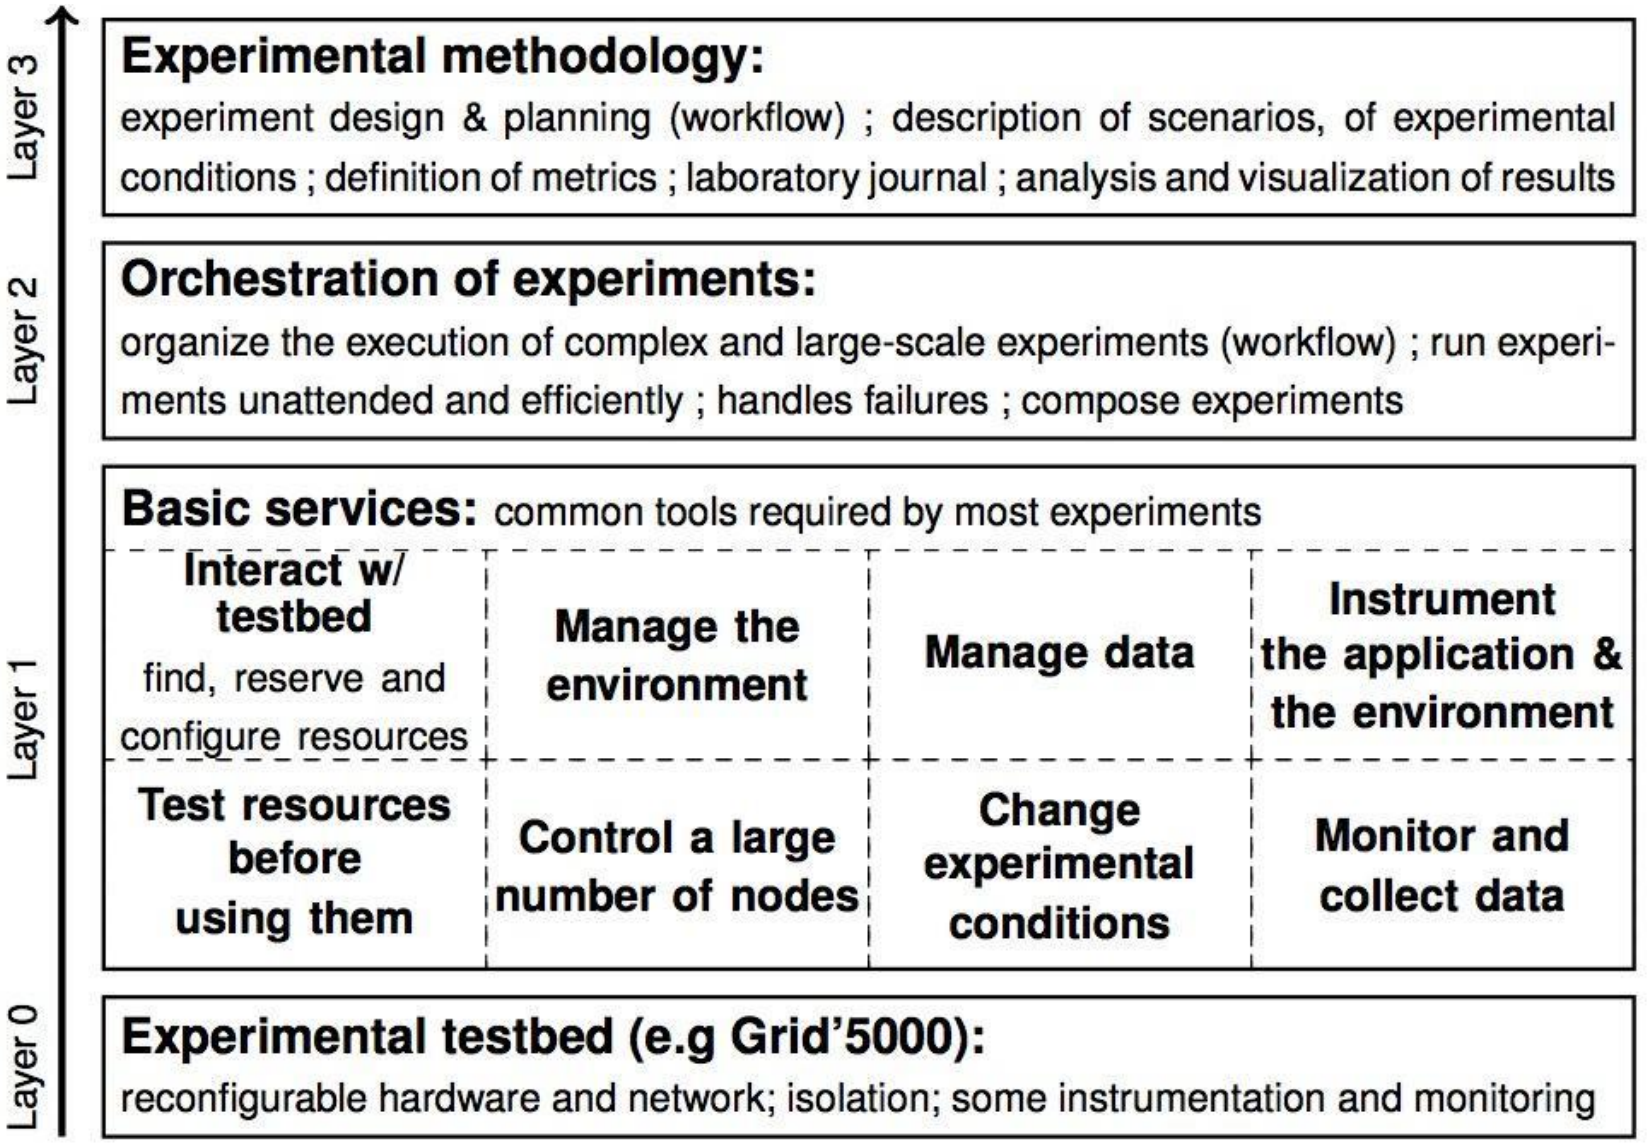
Layer 3
Experimental methodology:
experiment design & planning (workflow) ; description of scenarios, of experimental conditions ; definition of metrics ; laboratory journal ; analysis and visualization of results
Layer 2
Orchestration of experiments:
organize the execution of complex and large-scale experiments (workflow) ; run experiments unattended and efficiently ; handles failures ; compose experiments
Layer 1
Basic services: common tools required by most experiments
Interact w/ testbed
find, reserve and configure resources
Manage the environment
Manage data
Instrument the application & the environment
Test resources before using them
Control a large number of nodes
Change experimental conditions
Monitor and collect data
Layer 0
Experimental testbed (e.g Grid'5000):
reconfigurable hardware and network; isolation; some instrumentation and monitoring

Plans for SILECS: Testbed Services

- **Provide a unified framework that (really) meets all needs**
  - Make it easier for experimenters to move for one testbed to another
  - Make it easy to create simultaneous reservations on several testbeds (for cross-testbeds experiments)
  - Make it easy to extend SILECS with additional kinds of resources
  - Notes
    - SFA is probably not a sufficient solution here, even if SFA-compatibility is required for international collaboration (e.g. European federation)
    - It should be designed carefully to ensure we do not add just another not-so-useful abstraction layer
- **Factor testbed services**
  - Services that can exist at a higher level, e.g. open data service, for storage and preservation of experiments data (in collaboration with Open Data repositories such as OpenAIRE/Zenodo)
  - Services that are required to operate such infrastructures, but add no scientific value
    - users management, usage tracking

Services & Software Stack

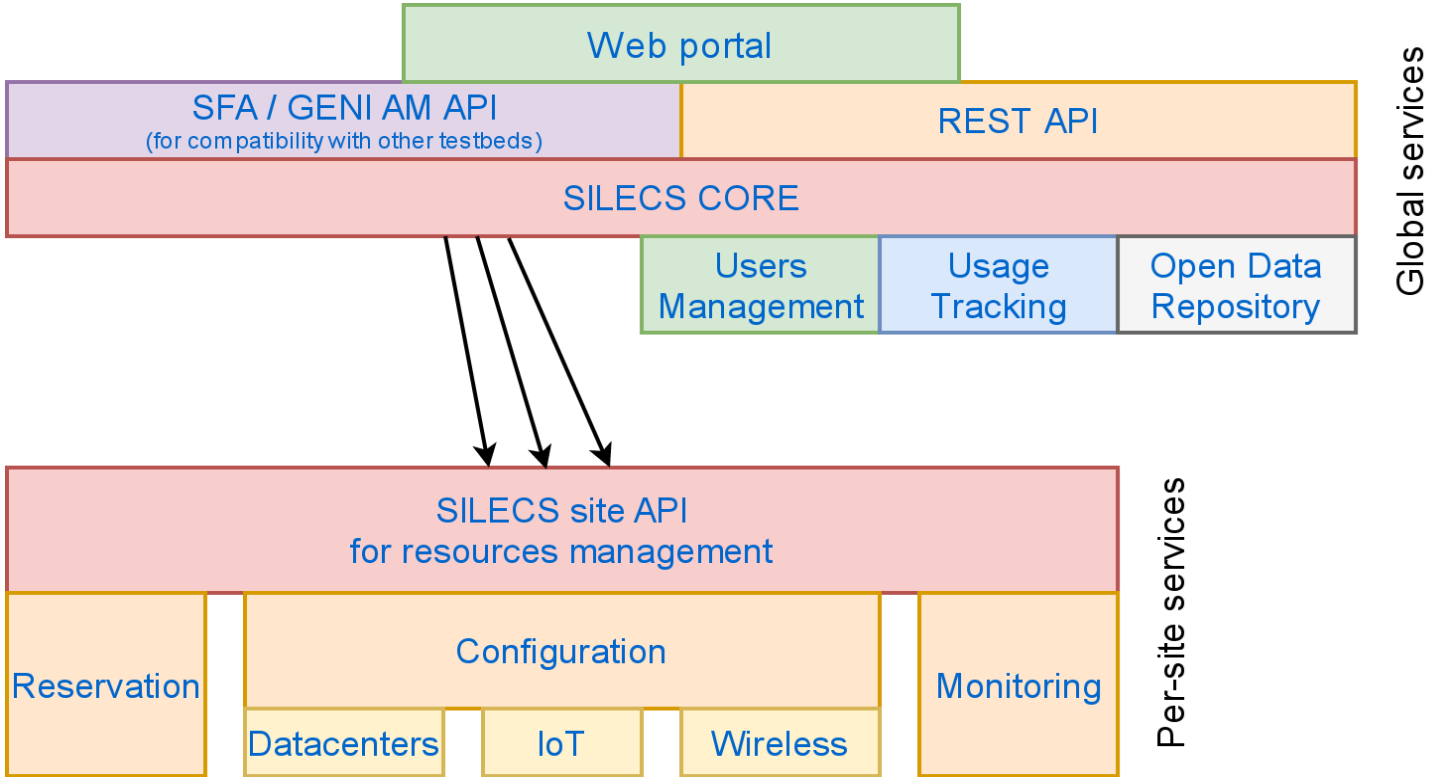

Built from already functional solutions

European Dimension (ESFRI)

Countries	FR	GR	CH	ES	CY	IT	DE	NL	LU	BE
										
Gov.										
										
Research	Inria cnrs SORBONNE UNIVERSITÉ UNIVERSITÉ DE STRASBOURG IMT	UNIVERSITY OF THESSALY CERTH CENTRE FOR RESEARCH & TECHNOLOGY HELLAS	MANDAT INTERNATIONAL ETH Zürich EPFL ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE	IMDEA networks Universidad Carlos III de Madrid UNIVERSIDAD DE GRANADA Universidad Zaragoza Universidad de Alcalá UPC Universidad del País Vasco / Euskal Herriko Unibertsitatea	UCLan CYPRUS	UNIVERSITÀ DEGLI STUDI DI MILANO ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA CRS4	JGU JOHANNES GUTENBERG UNIVERSITÄT MAINZ Fraunhofer	VU UNIVERSITY AMSTERDAM	uni.lu UNIVERSITÉ DU LUXEMBOURG	UNIVERSITEIT GENT
Industry	ERICSSON Hewlett Packard Enterprise Telefónica OHL Concesiones vodafone Teldat bintec-elmeg TTi AMETIC									
NRENs	RENATER CONNECTEUR DE SAVOIRS GÉANT									

Conclusions

- New infrastructure based on two existing instruments (FIT and Grid'5000)
- Design a software stack that will allow experiments mixing both kinds of resources while keeping reproducibility level high
- **Keep the aim of previous platforms** (their core scientific issues addressed)
  - Scalability issues, energy management, …
  - IoT, wireless networks, future Internet for SILECS/FIT
  - HPC, big data, clouds, virtualization, deep learning … for SILECS/Grid'5000
- **Address new challenges**
  - IoT and Clouds
  - New generation Cloud platforms and software stacks (Edge, FOG)
  - Data streaming applications
  - Big data management and analysis from sensors to the (distributed) cloud
  - Mobility
  - …

Thanks, any questions ?

http://www.silecs.net/

https://www.grid5000.fr/

https://fit-equipex.fr/