Grid'5000

FIT FUTURE INTERNET TESTING FACILITY

SLICES-RI

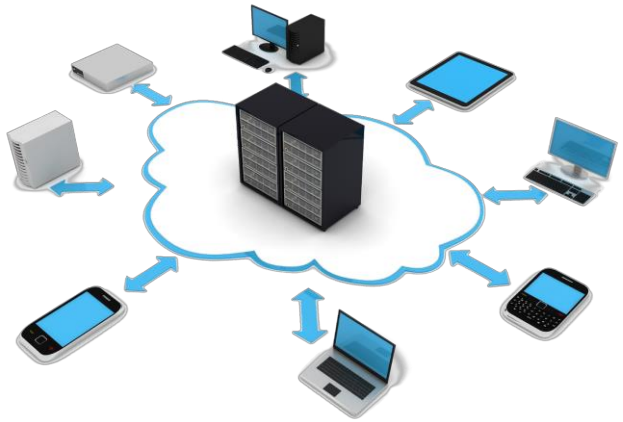

SILECS/SLICES

Super Infrastructure for
Large-Scale Experimental Computer Science

Christian Perez – LIP/Inria

Slides from F. Desprez – Inria/LIG & S. Fdida – Sorbonne University

INRIA, CNRS, RENATER, IMT, Sorbonne Université, Université Grenoble Alpes, Université Lille 1, Université Lorraine, Université Rennes 1, Université Strasbourg, Université fédérale de Toulouse, ENS Lyon, INSA Lyon, …

http://www.silecs.net http://slices-ri.eu

The Discipline of Computing: An Experimental Science

The reality of computer science

- Information
- Computers, networks, algorithms, programs, etc.

Studied objects are more and more complex

- hardware, programs, data, protocols, algorithms, networks

Example of multiple sources of complexity

- Processors have very nice features: caches, hyperthreading, multi-core, …
- Operating system impacts the performance
- The runtime environment plays a role (MPICH ≠ OPENMPI)
- Middleware have an impact
- Various parallel architectures that can be heterogeneous, hierarchical, distributed, dynamic

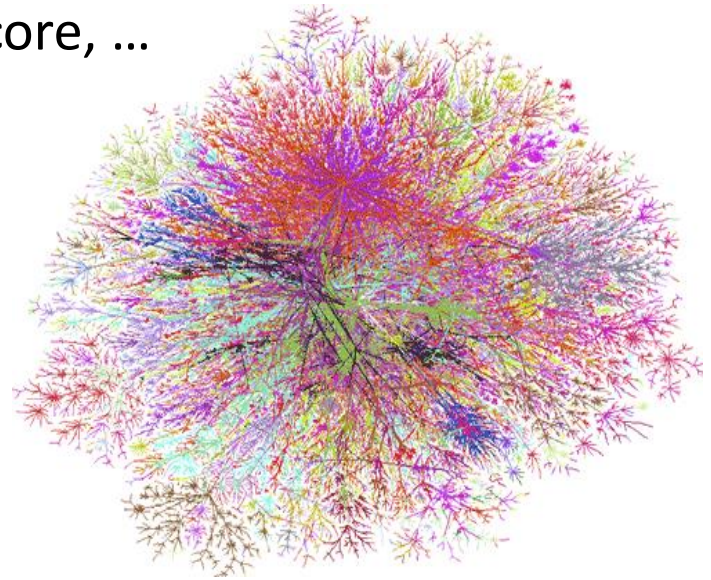

Convergence of Computation and Communication

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 democratisation

1999 - The Grid

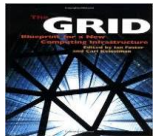

2002- Amazon Initial Compute/Storage services

2006 - Amazon EC2 (IaaS)

2015
Fog / Edge
Software Defined XXX

Transistor
1947

micro processor
1971

1838 - Telegraph

1876 - Telephone

1957 - satellite

1896 - Radio

1969 - ARPANET

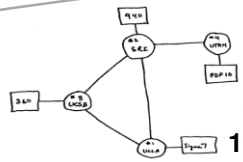

1985 - TCP/IP Adoption
1973 - Ethernet

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

ACM Artifact Review and Badging

https://www.acm.org/publications/policies/artifact-review-and-badging-current

SILECS/SLICES Motivation

- 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, …
  - Large variety of applications and large community
- 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)

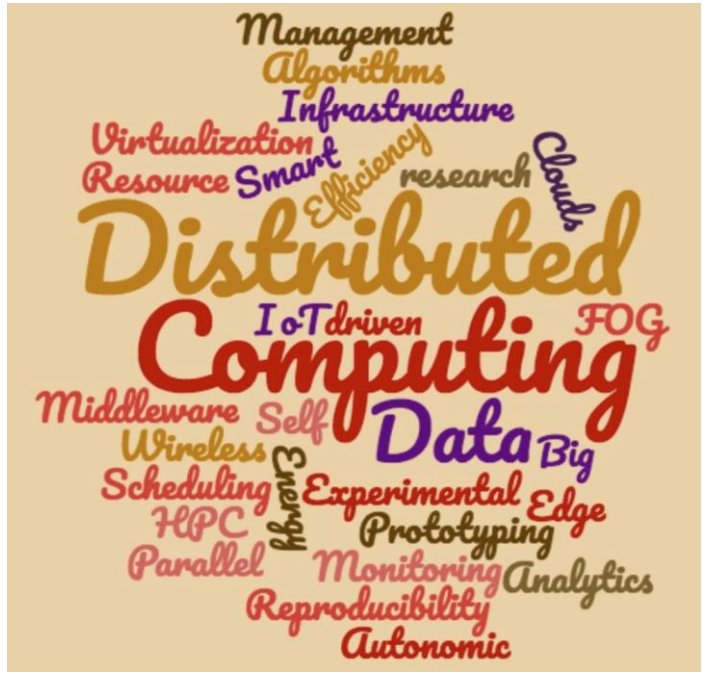

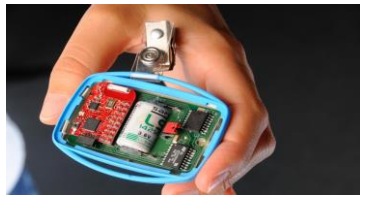

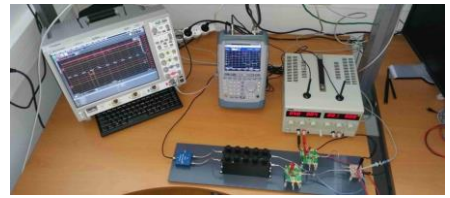

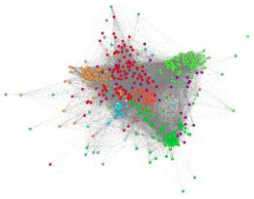

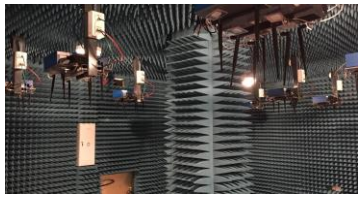

SILECS and SLICES

Need of specific platforms to experiment

- To measure how programs behave and not only of the results they produce
- To (dynamically) change the execution environment (up to generate real faults)
- Tier 0,1,2 only enable to execute « safe » programs

French level: Silecs

- Based upon two existing infrastructures: Grid'5000 (HPC/cloud) and FIT (wireless/IoT)
- On the f*euille de route nationale des Infrastructures de recherche* since 2018
  - https://www.enseignementsup-recherche.gouv.fr/pid25366/acces-thematique.html?theme=317&subtheme=318

Eurpean level: Slices

- In case of success, the first European the first Research Infrastructure in computer science

SLICES – ESFRI Call (Sept. 2020)

25 Participants from 15 countries

- Belgium
- Cyprus
- Finland
- **France (leader)**
- Germany
- Greece
- Hungary
- Italy
- Luxembourg
- The Netherlands
- Norway
- Poland
- Spain
- Sweden
- Switzerland

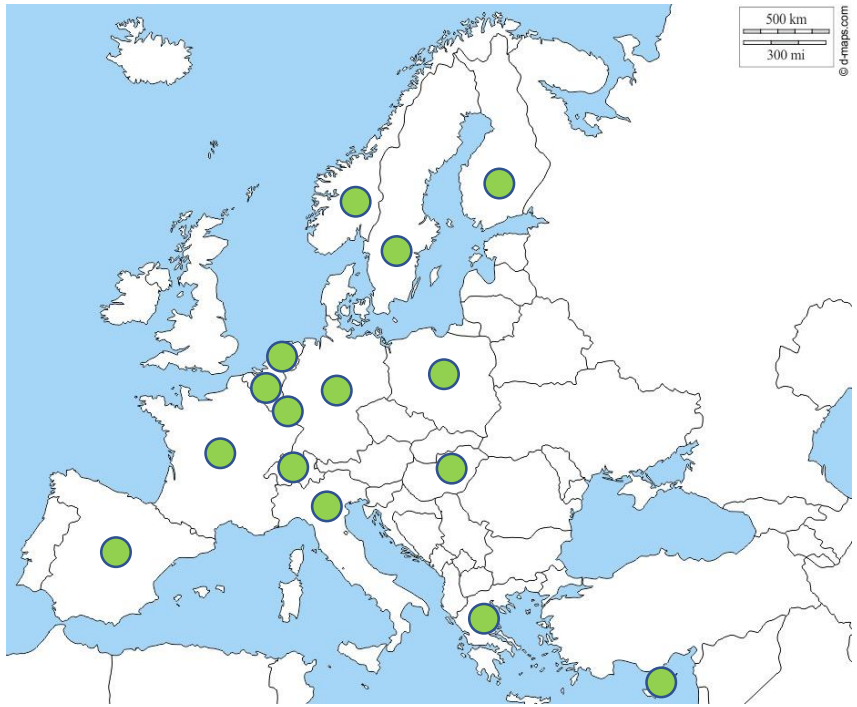

In cooperation with GIANT and national NRENs

Strong integration into the EOSC ecosystem

Timeline

Currently under evaluation

- Hearing @ Spring 2021

Design: 2017-2022*

Preparation: 2022-2025

Implementation: 2024-2028

Operation: 2024-2040

Termination: 2040-2042

Estimated total investment: 137 m€

* Supported by 2 projects started in 2020

- H2020 Slices Design Study
- H2020 Slices Starting Community

SILECS – PIA-3/EQUIPEX+ Call (June 2020)

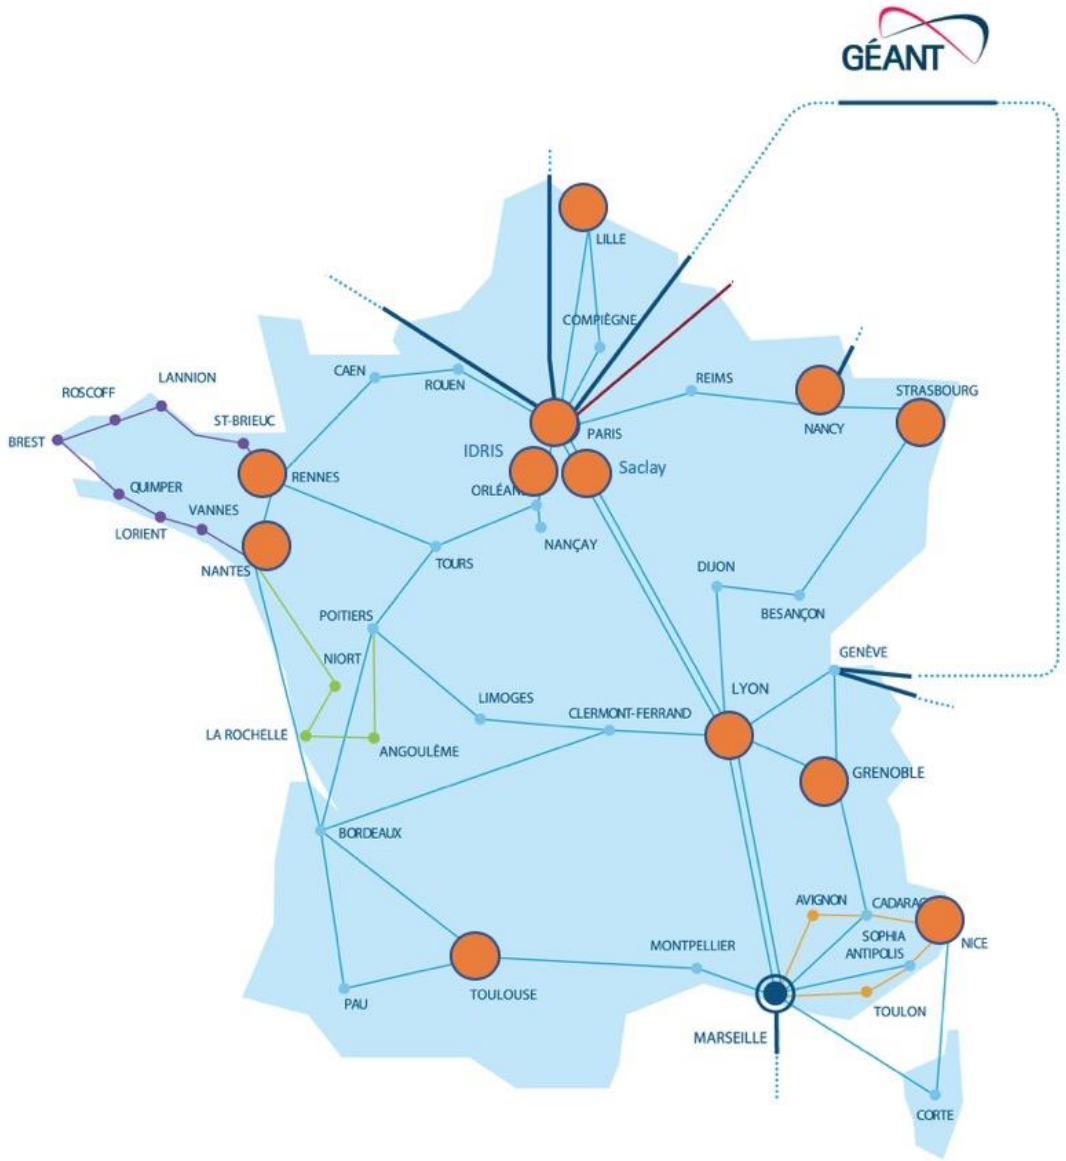

- **Core partners**
  - Inria
  - CNRS
  - IMT
  - Université fédérale de Toulouse
  - Université Strasbourg
  - Université Grenoble Alpes
  - Université de Lorraine
  - Sorbonne Université
  - Renater
  - *Eurecom*
  - *ENS Lyon*
  - *INSA de Lyon*

- **Other participants**
  - Université de Lille
  - Université de Rennes 1
  - Université de Lyon
  - Université de Nantes

Envisioned Architecture

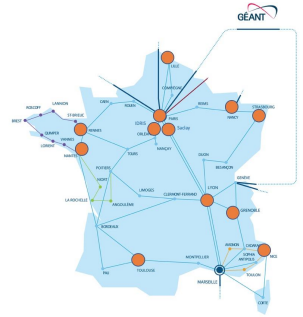

SILECS/GRID'5000

- **Testbed for research on distributed systems**
  - Born in 2003 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**
  - 8 sites, 36 clusters, 838 nodes, 15116 cores
  - 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

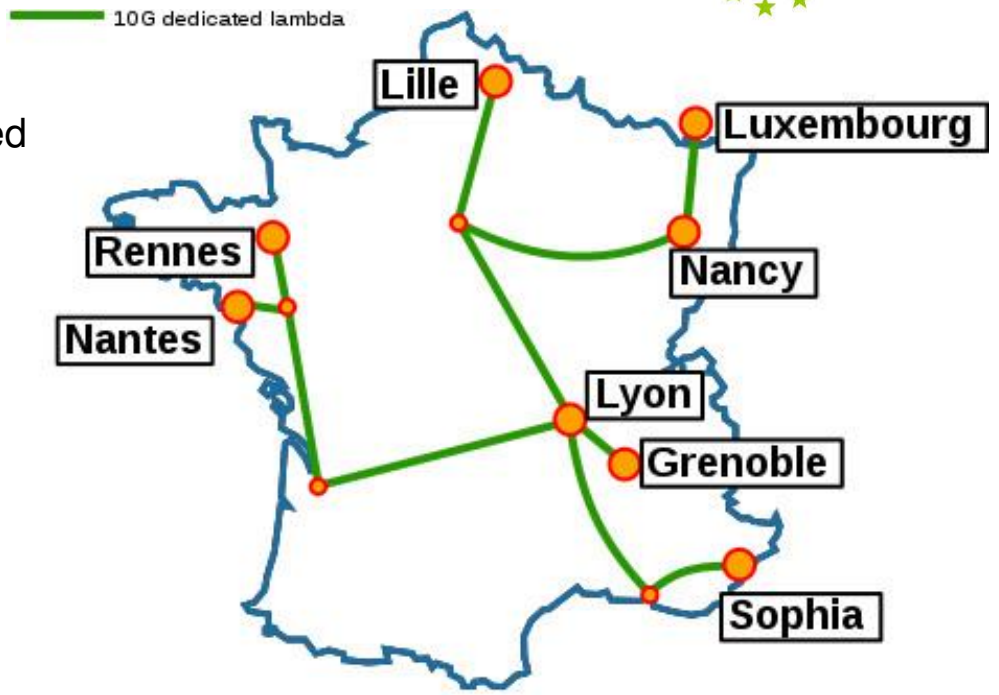

https://www.grid5000.fr/

SILECS/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

- **FIT-R2Lab:** WiFi mesh testbed (DIANA)
- **FIT-CorteXlab:** Cognitive Radio Testbed 40 Software Defined Radio Nodes (SOCRATE)

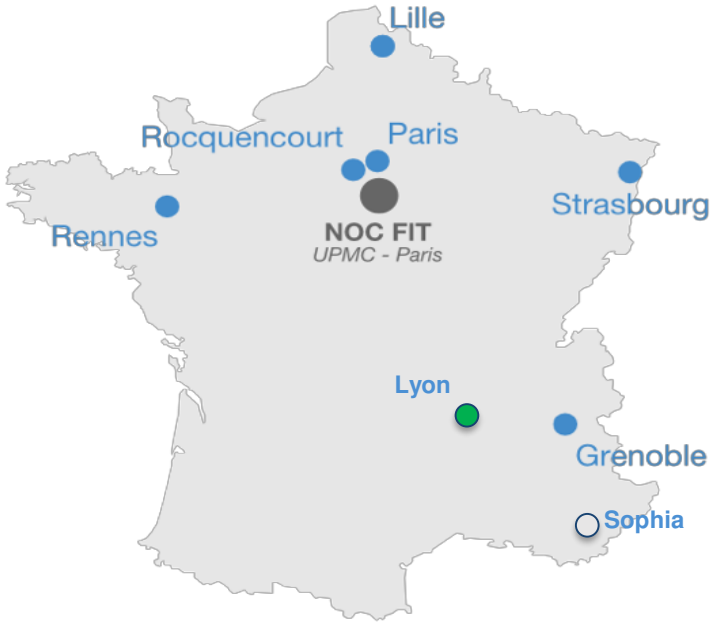

- **FIT-IoT-LAB**
  - 2700 wireless sensor nodes spread across six different sites in France
  - Nodes are either fixed or mobile and can be allocated in various topologies throughout all sites

https://www.iot-lab.info/hardware/

https://fit-equipex.fr/

SILECS: Data Center Portfolio

Targets

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

Hardware

- Servers: x86, ARM64, POWER, accelerators (GPU, FPGA), …
- AI dedicated servers
- Edge computing micro datacenters
- 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)

SILECS: 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), NB-IoT, 5G, BLE, Thread
- Wireless Sensor Network (IEEE 802.15.4),
- LoRa/LoRaWAN, …

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)

SILECS: Outdoor IOT testbed

- IoT is not limited to smart objects or indoor wireless sensors
  - smart building, industry 4.0, ….
- Smart cities need outdoor IoT solutions
  - Outdoor smart metering
  - Outdoor metering at the scale of a neighborhood (air, noise smart sensing, ….)
  - Citizens and local authorities are more and more interested by outdoor metering
- Controlled outdoor testbed
  - (Reproducible) polymorphic IoT: support of multiple IoT technologies (long, middle and short range IoT wireless solutions) at the same time on a large scale testbed
  - Agreement and support of local authorities
  - Deployment in Strasbourg city (500000 citizens, 384 km2)

An experiment outline

- Discovering resources from their description
- Reconfiguring the testbed to meet experimental needs
- Monitoring experiments, extracting and analyzing data
- Controlling experiments: API

Need to automatize/simplify the workflow and/or integrate it into higher level tool to enhance reproducibility

Plans for SILECS/SLICES: 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/SLICES with additional kinds of resources

- **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 and EOSC
  - Services that are required to operate such infrastructures
    - Users management, usage tracking, etc.

Services & Software Stack

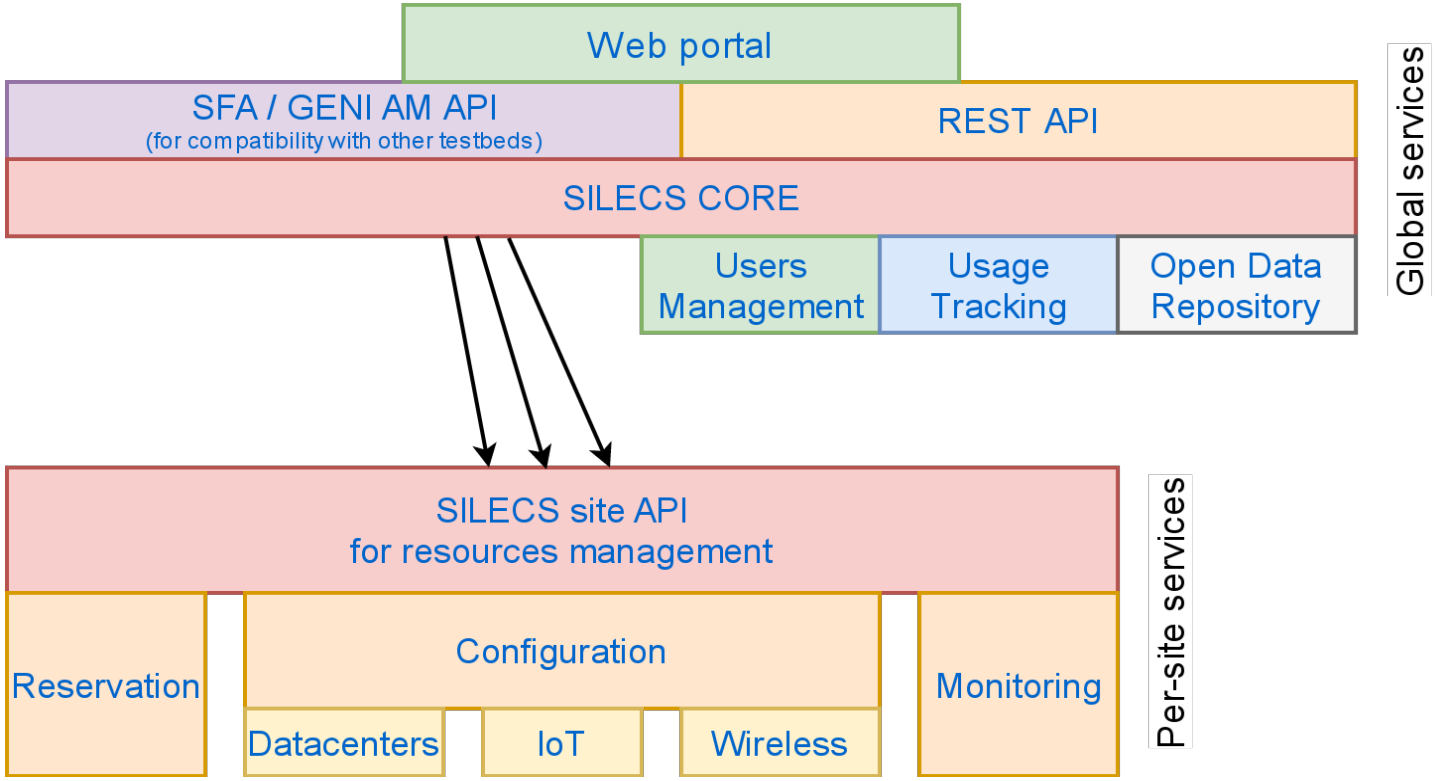

Built from already functional solutions

Exchanges with the community

- **JCAD 2018, 2019, and 2020**

- **TILECS Workshop**
  - 2019, July 3-4, Grenoble, LIG/IMAG
  - 101 attendees (academics and some from the industry)
  - **https://www.silecs.net/tilecs-2019/**

- **Silecs – Request for input (closed)**
  - 1/2 page(s) document describing which kind of experiment you would like to perform in the next 4 years and what will be you dream infrastructure (hardware/software/services)
  - Analysis to be provided soon

- **Slices – Request for input (open)**
  - Online survey on Research Infrastructure Needs and Requirements (15 questions)
  - **https://survey.iotlab.eu/index.php/326467?lang=en**

Conclusions

- **SLICES**: new infrastructure for experimental computer science and future services in Europe
- **SILECS**: new infrastructure in France based on two existing instruments (FIT and Grid'5000)
- **Big challenges !**
  - Design a software stack that will allow experiments mixing both kinds of resources while keeping reproducibility level high
  - Keep the existing infrastructures up while designing and deploying the new one
- **Keep the aim of previous platforms** (their core scientific issues addressed)
  - Scalability issues, energy management, …
  - IoT, wireless networks, future Internet
  - HPC, big data, clouds, virtualization, deep learning, ...
- **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
  - Next generation wireless
  - …
- **Next steps**
  - Waiting for results for PIA-3EQUIPEX+ and ESFRI Slices (hearing in Spring 2021)

Thanks, any questions?

http://www.slices-ri.eu

https://www.silecs.net/

https://www.grid5000.fr/

https://fit-equipex.fr/