SLICES

European Scientific Large-Scale Infrastructure for Computing/Communication Experimental Studies

SLICES at INDIS Workshop SC23

November 2023

Dr. Paola Grosso – University of Amsterdam

Outline

- SLICES Overview
- Research prioritization and Post 5G experimental facilities and testbed
- Experimental Research Reproducibility in SLICES
- SLICES Data Management Infrastructure

From mid-Scale (~100M€) to Large-Scale (~B€)

ESFRI

MAKING SCIENCE HAPPEN

A new ambition for Research Infrastructures in the European Research Area

The European ESFRI framework

European Strategy Forum on Research Infrastructures

Supporting a scientific methodology

http://www.esfri.eu/

slicesRI

Third generation Mid-Scale Test Platforms

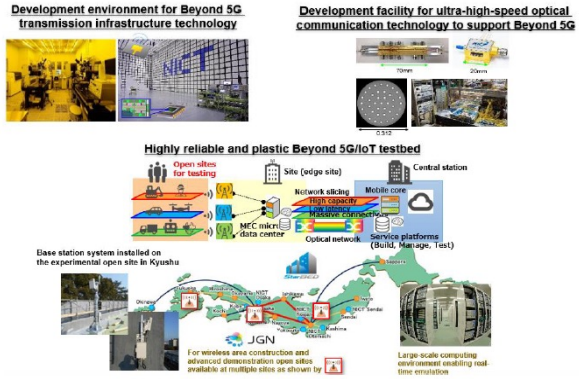

USA NSF PAWR (Platforms for Advanced Wireless Research): NSF + Industry, 100M€, 2017-2022

NSF Fabric: NSF, 20 M€, 2019-2023

Colosseum: NSF-DARPA, 20+7,5M$, 2017-2025.

BRIDGES: NSF, 2.5M€, 2020-2023

EU Horizon Europe
ICT 17-19-52, 2018-2022, 205 M€
SNS Stream C, first call, 2022-2025, 25M€

Japan NICT R&D
Shared Open Platform
200 M$

China CENI
Chinese Experimental National Infrastructure
2018-2022
190 M€

slicesRI

SLICES for research on Digital Infrastructures

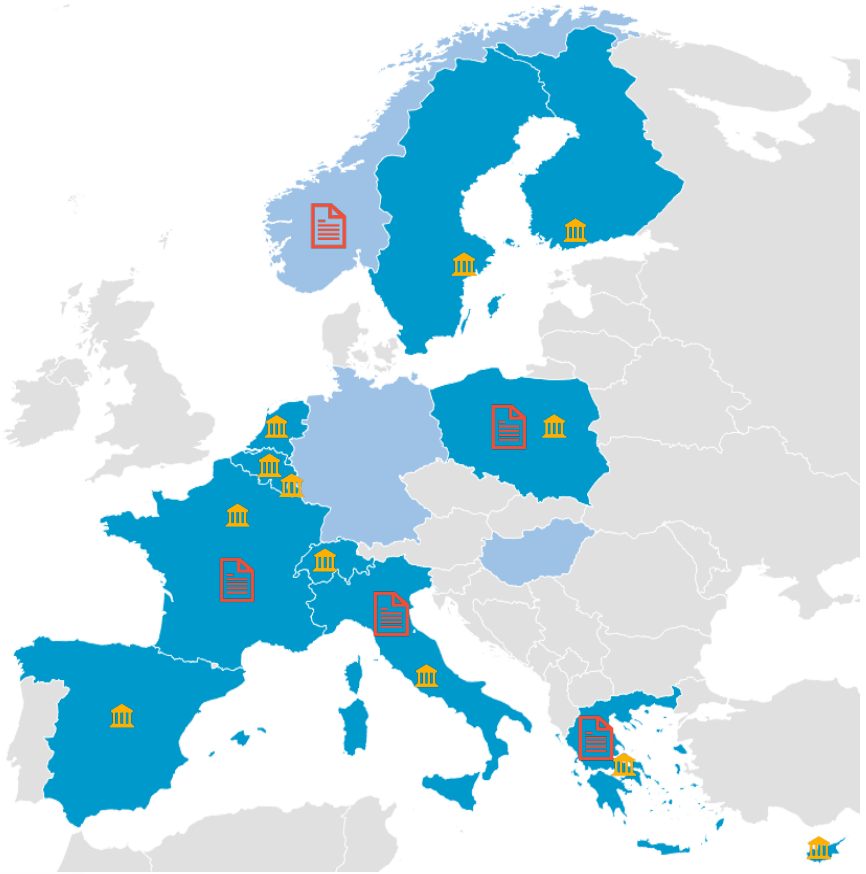

Initiated in 2017, **25 partners** from 15 countries:

- **12 political support** from National Ministries
- included in **5 national roadmaps**

SLICES will enable **scientific excellence and breakthrough** and will **foster innovation in the ICT domain**, strengthening the **impact of European research**, while contributing to European agenda to address **societal challenges**, and in particular, the twin transition to a sustainable and digital economy.

slicesRI

Current status of the partnership

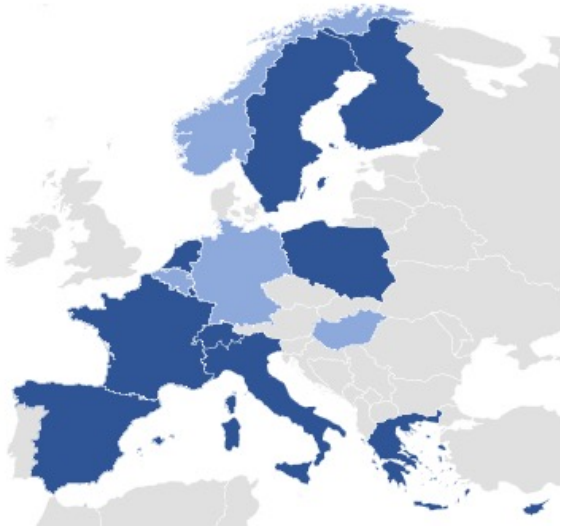

SLICES
ESFRI successful application – 2020

slices RI

Countries	Government	Research and Academia		Industry	Clusters, networks and others	NRENs	Worldwide support
	National support	*Partners*	*Support*				
Core partners							
	Local support confirmed						

Countries (15)

- France
- Belgium
- Cyprus
- Finland
- Germany
- Greece
- Hungary
- Italy
- Luxemburg
- Netherlands
- Norway
- Poland
- Spain
- Sweden
- Switzerland

Worldwide Cooperation

NSF

6G SNS IA

Google Amazon Facebook Apple Microsoft

ONAP
OPEN NETWORK AUTOMATION PLATFORM

O-RAN
ALLIANCE

ONF
Open Networking Foundation

OPEN AIR INTERFACE

slicesRI

SLICES-PP (2022-2025): Consortium members

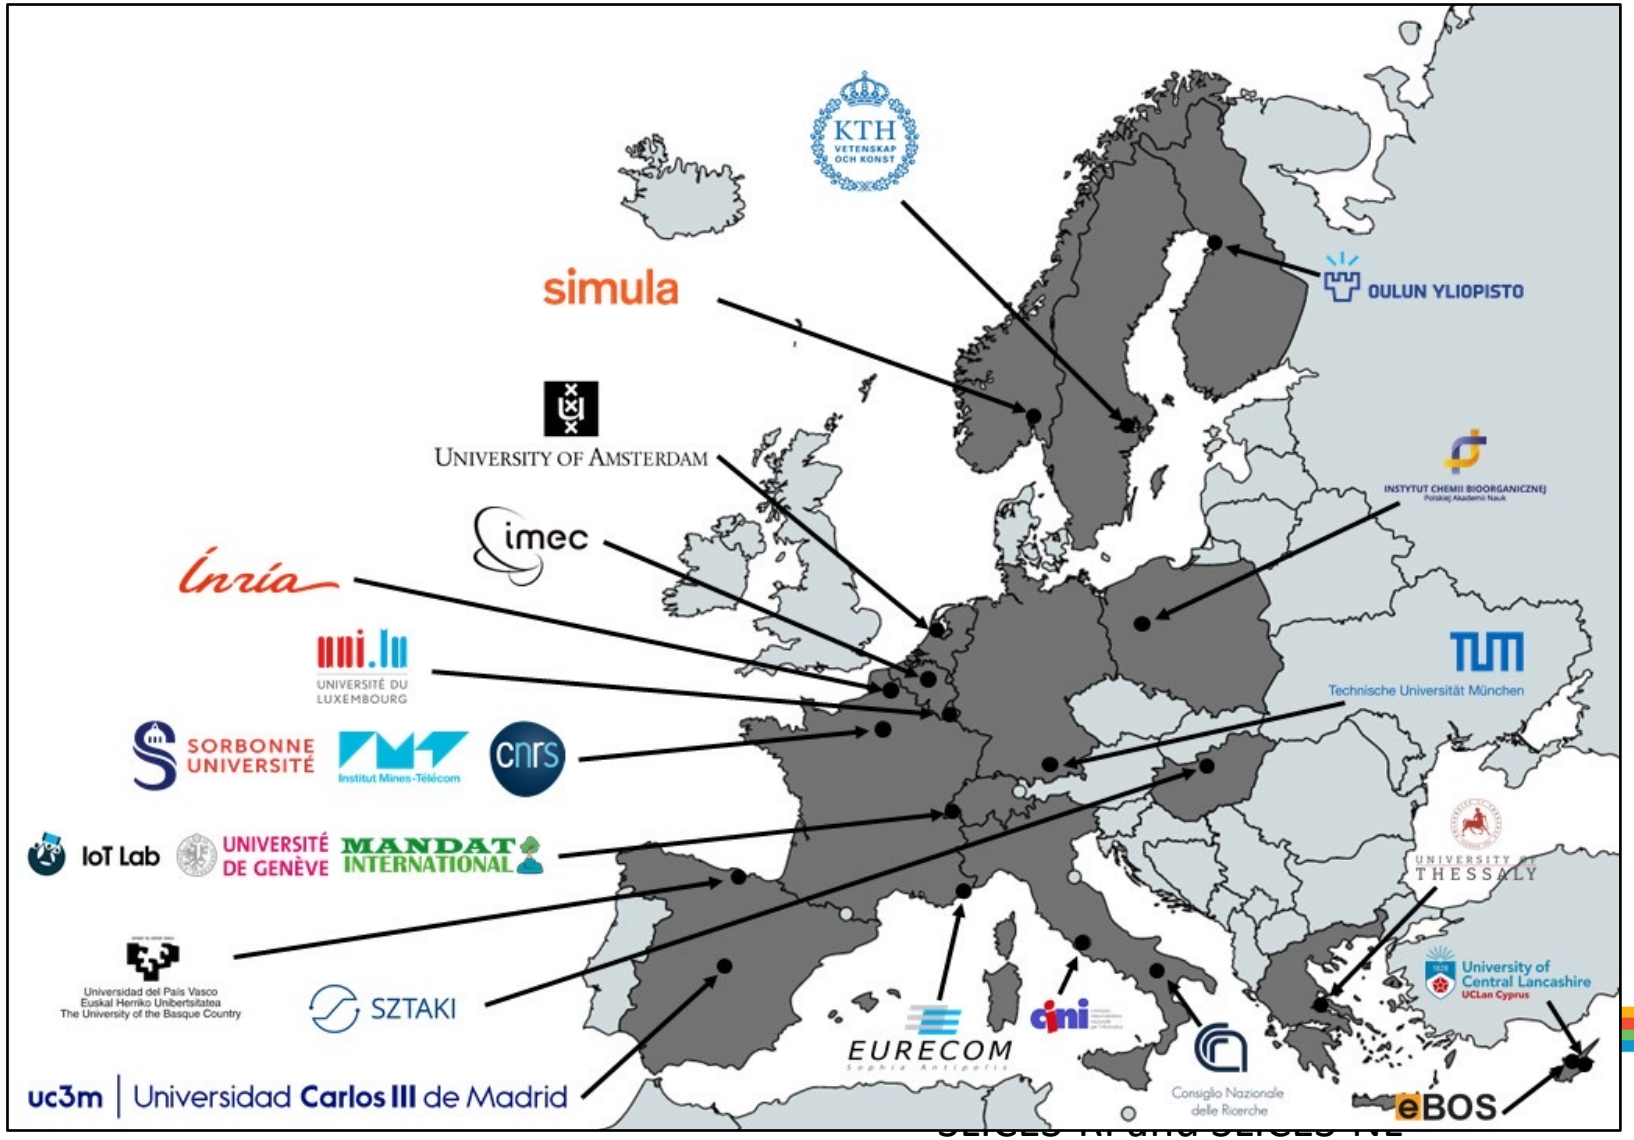

25 Partners from 15 countries

- INRIA, FR
- Sorbonne University (SU), FR
- **Univ of Amsterdam (uvA), NL**
- Univ of Thessaly (UTH), GR
- CNR, IT
- PSNC, PL
- Mandat International (MI), CH
- IoTLAB, FR
- UC3M, ES
- IMEC, BE
- UCLan, CY
- EURECOM, FR
- SZTAKI, HU
- CINI, IT
- CNIT, IT
- Univ Luxemburg, LU
- TUM, DE
- EHU, ES
- KTH, SE
- Univ Oulun, FI
- EBOS, CY
- SIMULA, NO
- IMT, FR
- Univ Geneve, CH

Testbeds running by SLICES-RI Partners

- OneLab: Cloud Infrastructure for Researchers (LIP6, Sorbonne University)
- imec testbed for networking, cloud, AI and IoT research (Ghent Uni)
- IoT Lab (Mandat International, CH)
- 5TONIC Lab (Uni Carlos III of Madrid)
- LeonR&Do Lab (COSMOTE, GR)
- NITOS testbed 5G (University of Thessaly)
- Open5G Lab, SOPHIA-NODE: Beyond-5G cloud-native network (INRIA)
- SN4I Lab Smart Network for Industry 4.0 (Uni Basque Country)

Prioritisation of research topics

What's the methodology behind it?

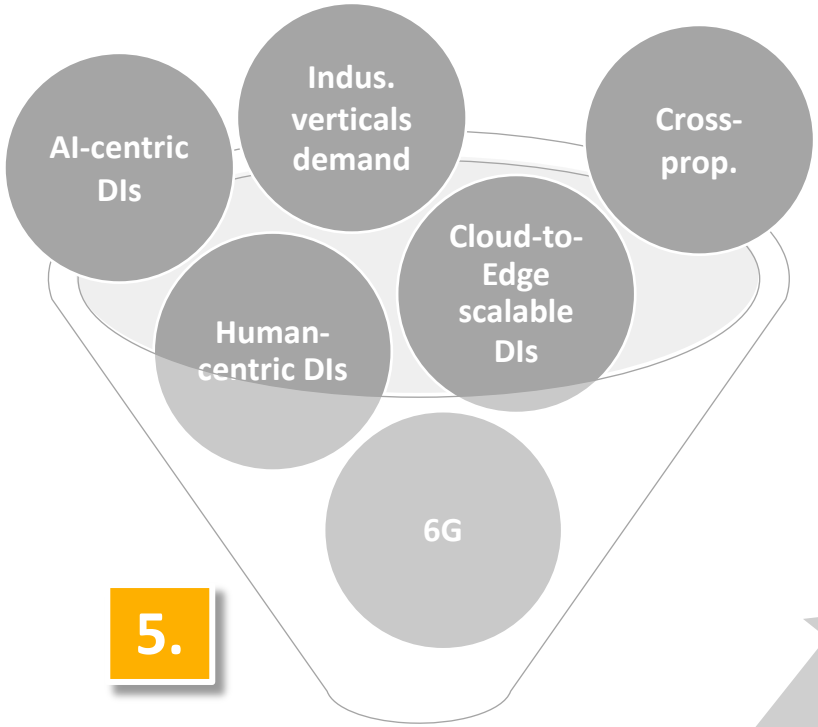

5.

- New waveforms, higher frequencies up to THz.
- Spectrum and wireless management.
- Integrated sensing and communication.
- Heterogeneous radio management.

ADVANCED WIRELESS NETWORKING

- Advanced protocols and architectures (virtualization, softwarization, programmability).
- AI applied to infrastructure operation and optimization.
- Generation of data to train algorithms.
- Distribution of intelligence into (and beyond) the Edge of the network.

SMART INFRASTRUCTURE OPERATION AND MANAGEMENT

- Fog/Edge/cloud hyper converged infras
- Software component deployment.
- Distributed resource management & microservices.
- Geo-distributed data management.
- Federated deep learning.
- Datacentres infras for distributed systems, appli. and software stacks.

DESIGN & VALIDATION OF NEW DIS AND HYPER-CONVERGED INFRAS

- New challenges arising from the verticals and the ubiquitous networks.
- Interoperability, composable infrastructure services on-demand (RI as a Service).
- Seamless user experiences across technologies and domains.

ADVANCED FUNCTIONALITIES

ENERGY EFFICIENCY AND CARBON FOOTPRINT

SECURITY AND PRIVACY

Breaking down in priority research topics

Simultaneous but progressive exploration of research topics

slicesRI

SLICES-RI PoC Blueprint – post5G Cloud-Edge

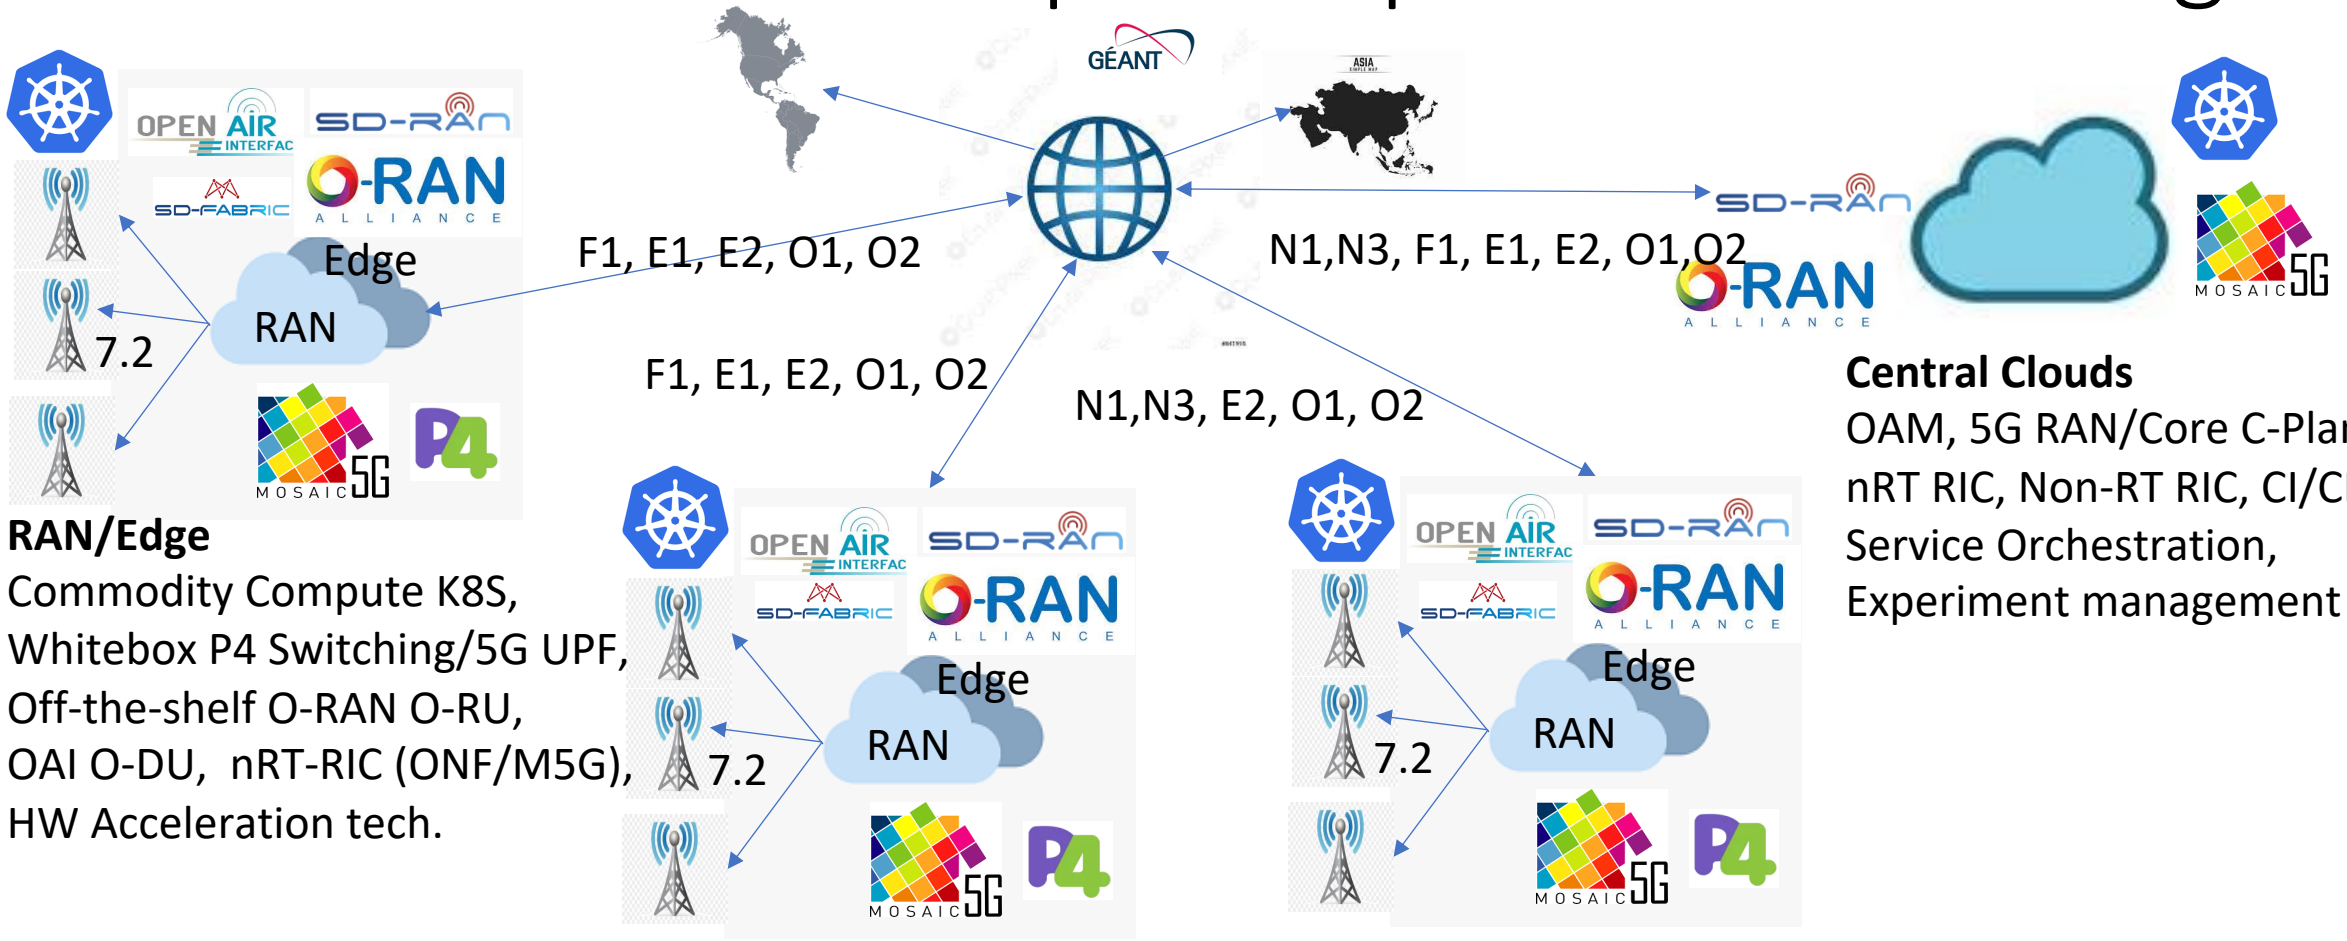

RAN/Edge
Commodity Compute K8S,
Whitebox P4 Switching/5G UPF,
Off-the-shelf O-RAN O-RU,
OAI O-DU, nRT-RIC (ONF/M5G),
HW Acceleration tech.

Central Clouds
OAM, 5G RAN/Core C-Plane,
nRT RIC, Non-RT RIC, CI/CD,
Service Orchestration,
Experiment management

slicesRI

SLICES Node preliminary blueprint – User Plane

5G Experiments Case

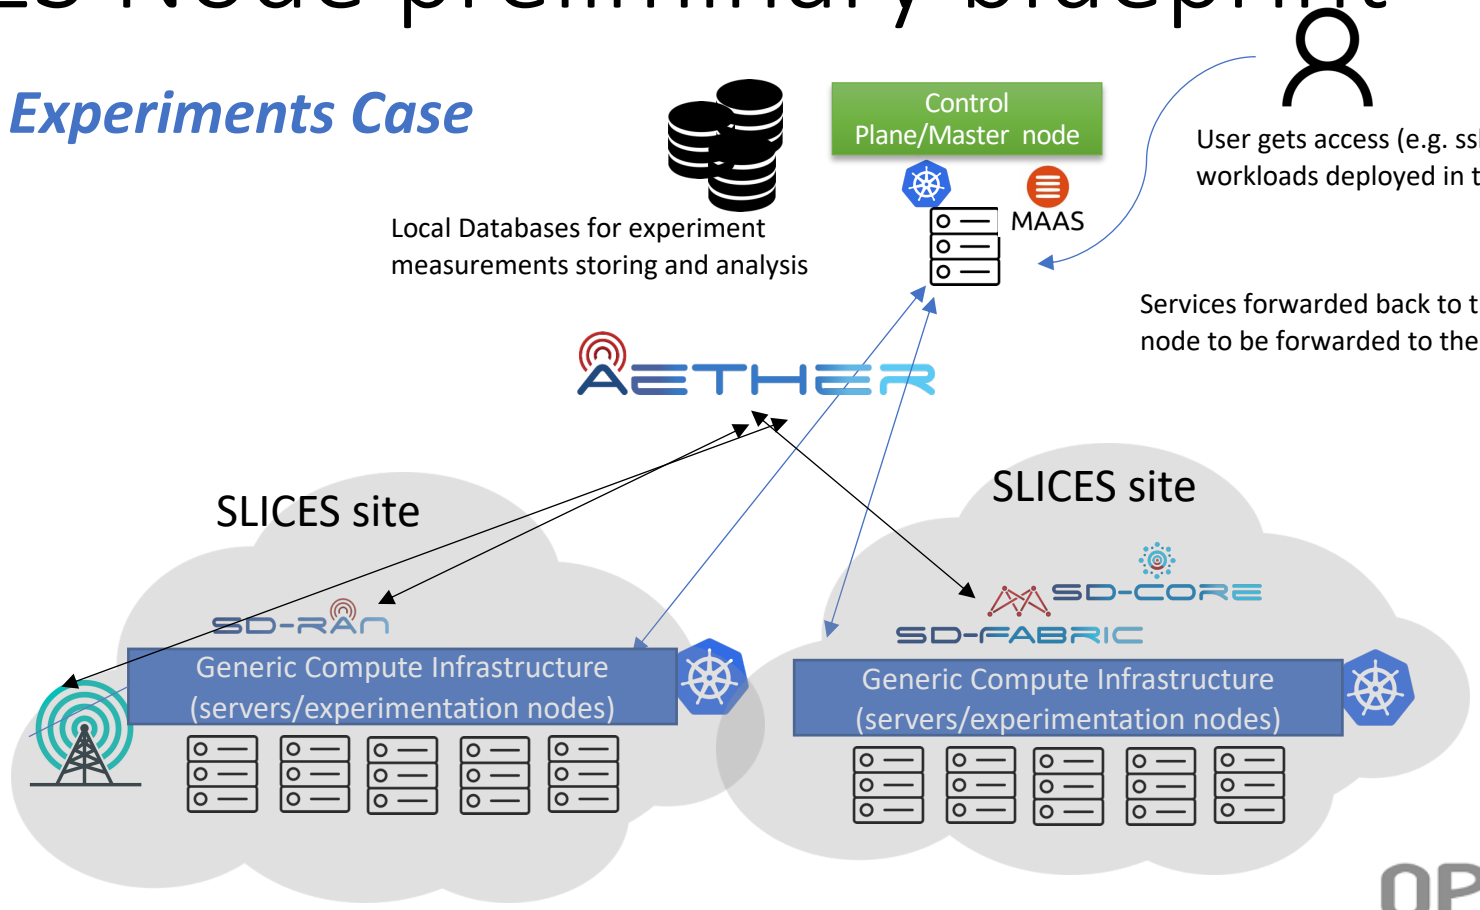

Can be a BS or an OAI USRP reachable from the respective BS-function (e.g. DU/RU) on the compute infrastructure

OAI Core Network (AMF, SMF, UDM, UDR, NRF, NSSF) Together with ONF UPF combined with P4 Switches

Blueprint User Perspective/Services

User Perspective - Experiments

Joint use of Post5G infrastructure services and SLICES-RI HPC resources

- **Example 1: Real-time Digital twin of radio network**
  - GPU farms can be used as real-time 3D radio emulators. When interconnected with radio and core network infrastructure can make a digital twin of a deployed network
  - This requires tight interconnection between radio processing infrastructure and the GPU farm but can be used to perform experiments not possible on the real network (large number of terminals).
  - Novel aspect, joint radio and digital twin. This requires proximity of HPC and real radio infrastructure

- **Example 2: Code analysis and bug fixing**
  - Protocol implementations are bug-ridden. In the CI/CD Type 3 experiment, developers of OSS networking software can make use of SLICES-RI GPU farms for code analysis and bug fixing.
  - Today in projects like OAI, CI makes use of “simple” tools like `cppcheck` to analyze community contributions. Use of AI/ML tools will take this to another level.

Roadmap:
Post 5G Testbed and Experimental Facilities

- ***July 2023 - Initial PoC***
  - IEEE HPSR Tutorial (USA)
  - EUCNC demo
  - SLICES-SC Summer School tutorial
- ***June 2023 - December 23 lessons / Deployment . Consolidation . Lessons learned***
  - Buildup of initial SLICES-RI post-5G sites (Targeting 6 countries – 10 physical sites)
  - Blueprint will provide input for planning new sites
  - Alignment with O-RAN NGRG platform activities and SNS Streams C/D
  - Alignment with other International activities (OpenRANGym, Japan, Brazil, 6G hubs in Germany)
- ***January+ 2024 – SLICES-RI Pre-operation***
  - Development of required interfaces for SLICES-RI (portal, central cloud services, contribution to API development)
  - CD activities

slices RI

Experimental Research Reproducibility in SLICES

Motivation for Reproducibility

Problems with reproducibility

- Two workshops at SIGCOMM conference dedicated to reproducible research:
  - SIGCOMM'03: MoMeTools workshop
  - SIGCOMM'17: Reproducibility workshop
  - Problems remained the same over 14 years

- Best solution so far…
- Artifact Evaluation Committees & Reproducibility Badges
- Problems:
  - High effort
  - Potentially low robustness (CCR Apr. '2021 [1])

ACM's badges awarded by the Artifact Evaluation Committee

[1] N. Zilberman, "An Artifact Evaluation of NDP", Comput. Commun. Rev., vol. 50, no. 2, pp. 32–36, 2020

What is Reproducibility?

3-stage process according to ACM [2]:

1. **Repeatability:** *Same* team executes experiment using *same* setup
2. **Reproducibility:** *Different* team executes experiment using *same* setup
3. **Replicability:** *Different* team executes experiment using *different* setup

A testbed-driven approach targets the experimental setup:

- Focus on **repeatability** and **reproducibility**
- **Replicability** requires additional effort by others

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

Experimental Research Reproducibility as a Service

SLICES to support experiments reproducibility to comply with Open Science

- Focus on **repeatability** and **reproducibility** with the future support of **replicability**

Robust, reproducible experiments

- Documenting all relevant parameters and environment for experiments
- Automate the documentation of experiments

➢ Well-structured experiment workflow may serve as documentation

Benefits for research community

- Reduce amount of work for experimenters to create reproducible experiments
- Reduce amount of work for other researchers to recreate and re-run experiments
- Make reproducibility an integral part of experiment design

➢ Automate entire experiment (setup, execution, evaluation)

Experimental research stages

- Experiment Planning
- Experiment setup, Equipment configuration
- Load (test) data
- Execute workflow
- Collect data
- **Evaluate and re-run experiment if needed**
- Process/analyse data
- Produce report
- Archive/publish data

slicesRI

Testbed-driven Experiments

The plain orchestrating service (pos) [3], a framework for reproducible experiments:

1. A testbed management system
2. A well-defined experiment workflow

Achieving Repeatability

- Automation & Linux Live Images
  - Researchers **must** automate configuration
  - No residual state between reboots
- ➢ Experiments become **repeatable**

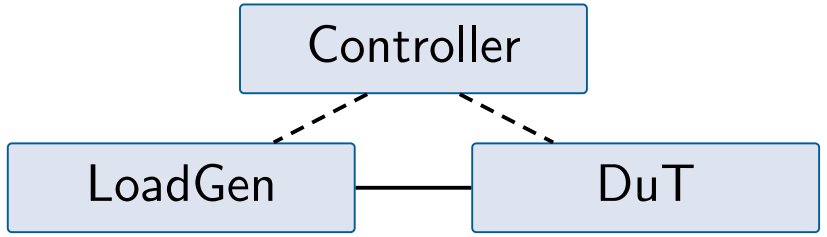

Minimal Experiment Topology

Achieving Reproducibility

- Providing access to experiment infrastructure
- Other researchers can easily (re-)run experiment
- ➢ Experiments become **reproducible**

[3] S. Gallenmüller, D. Scholz, H. Stubbe and G. Carle, "The pos Framework: a Methodology and Toolchain for Reproducible Network Experiments," in ACM CoNEXT'21, https://dl.acm.org/doi/10.1145/3485983.3494841

Plain Orchestration Service (pos) by Technical University Munich

- The plain orchestrating service (pos) provides two components:
  - Testbed controller and Experiment workflow
- The testbed controller takes care of the allocation and management of experimental resources
  - It provides bare-metal access to the experiment nodes
  - Images for the experiment nodes are provided in the form of live Linux images
- Using live images for experiments has two benefits:
  - First, rebooting an experiment node helps reset the system to a well-defined state.
  - Second, testbed users are aware of the non-permanence of their configuration, gently pushing users towards documenting and automating experiment configuration.

SLICES to provide the Robust Data Infrastructure for Experiment/Data Driven Research

- **Experimental data are big, distributed, domain specific, serving specific communities**
  - **Require effective models and infrastructure services for Research Data Management and secure data sharing**
- Support the whole data lifecycle
  - Connected to research/experiment lifecycle or workflow
- Distributed data storage and experimental data(set) repositories
  - Supporting recognized data interoperability standards (data formats and metadata)
  - Eventually certified: RDA endorsed Maturity and certification practice
  - **Interoperability and integration with EOSC as Federated data infrastructure**
- Data management and data curation and quality assurance
  - FAIR data principles and SLICES metadata profiles (interoperable with EOSC)
- Linked data and data discovery using semantic search and knowledge graph
  - PID (Persistent IDentifier) and FDO (FAIR Digital Object) infrastructure (interoperable with EOSC)
- (Trusted) Data exchange and secure transfer protocols

slicesRI

Different Types of Data for Different Experimental Studies

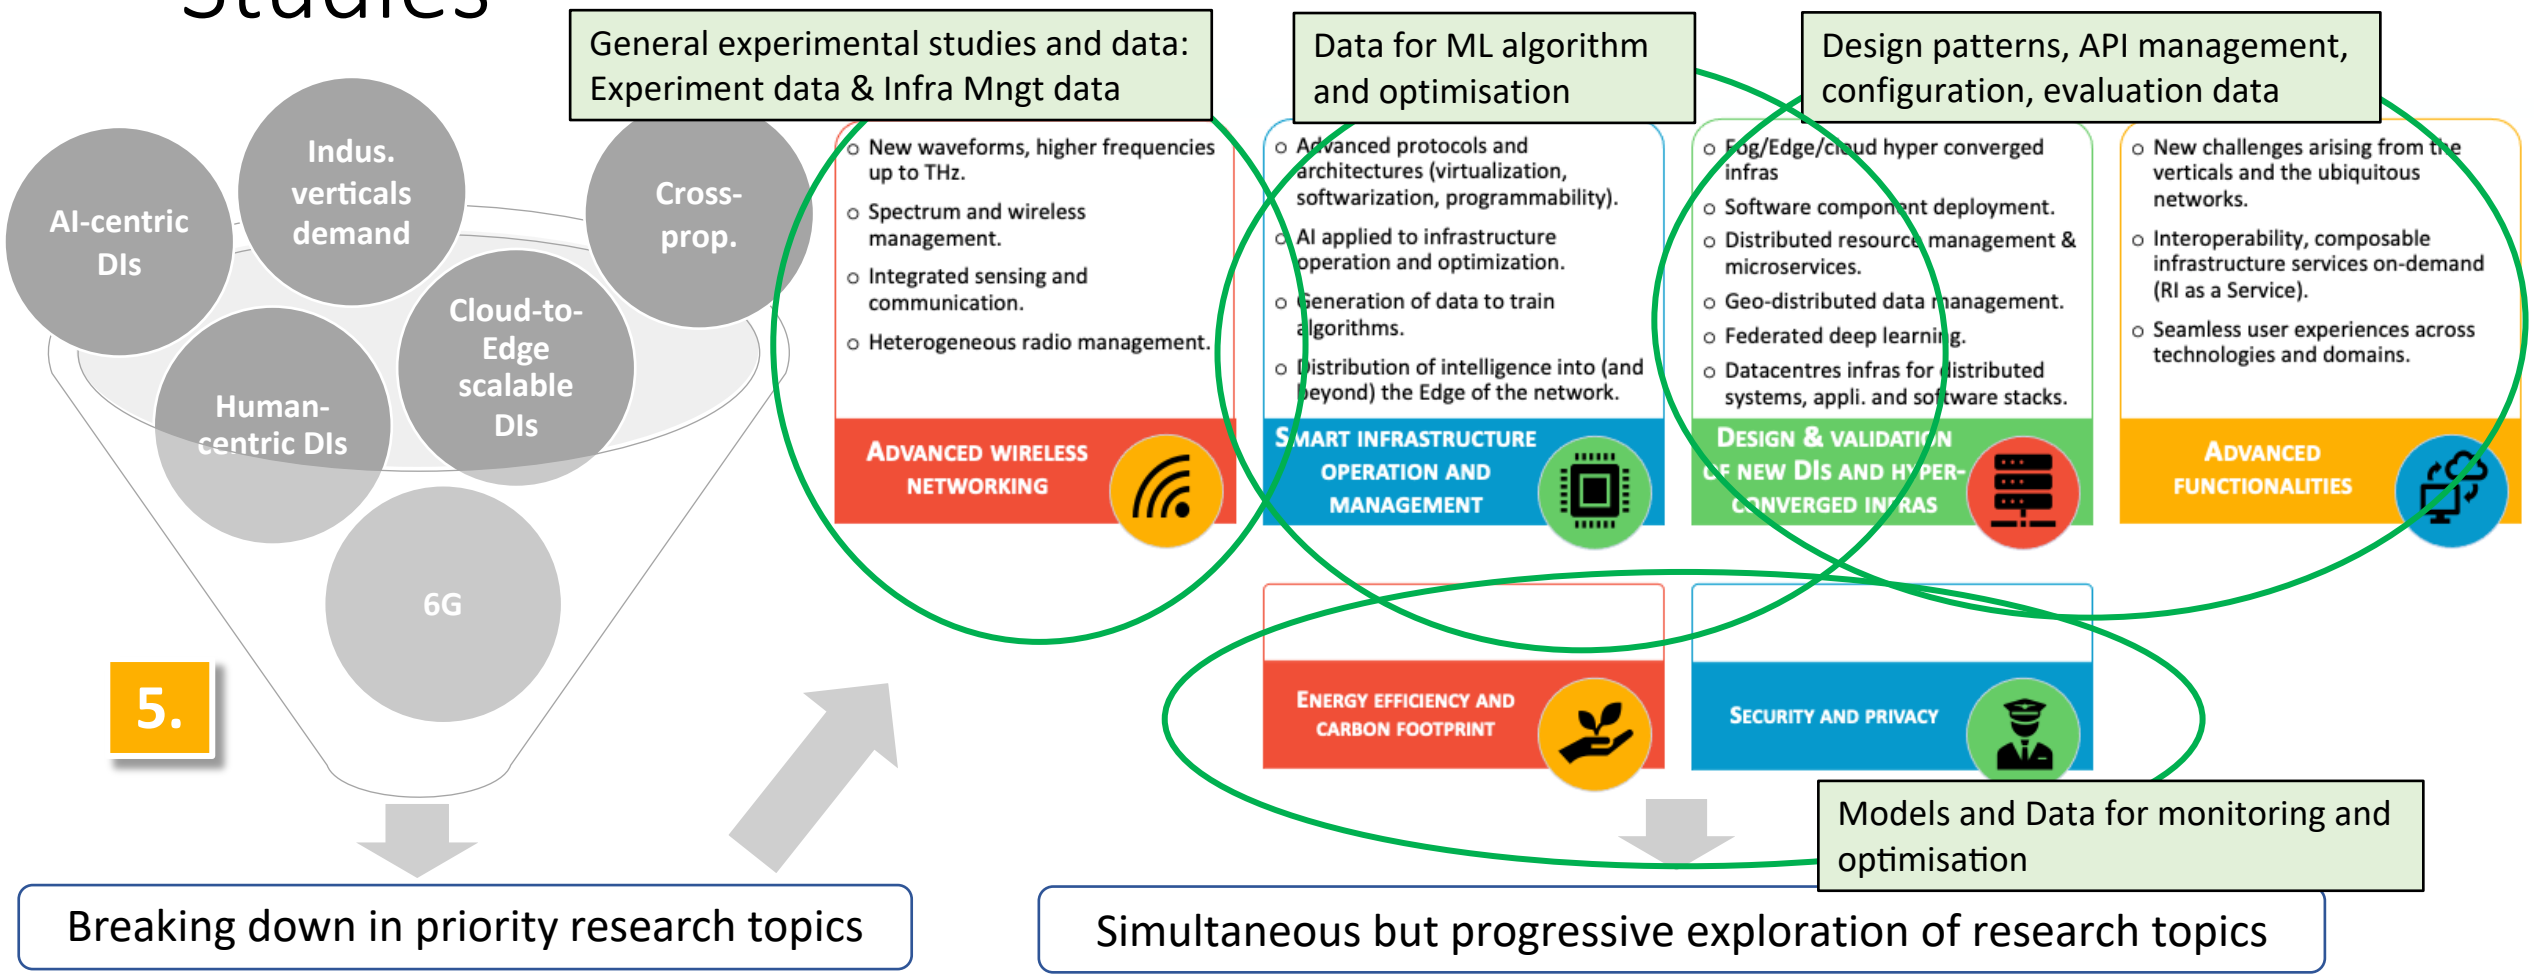

slicesRI

SLICES Experimental Data Lifecycle Model and Dataflow

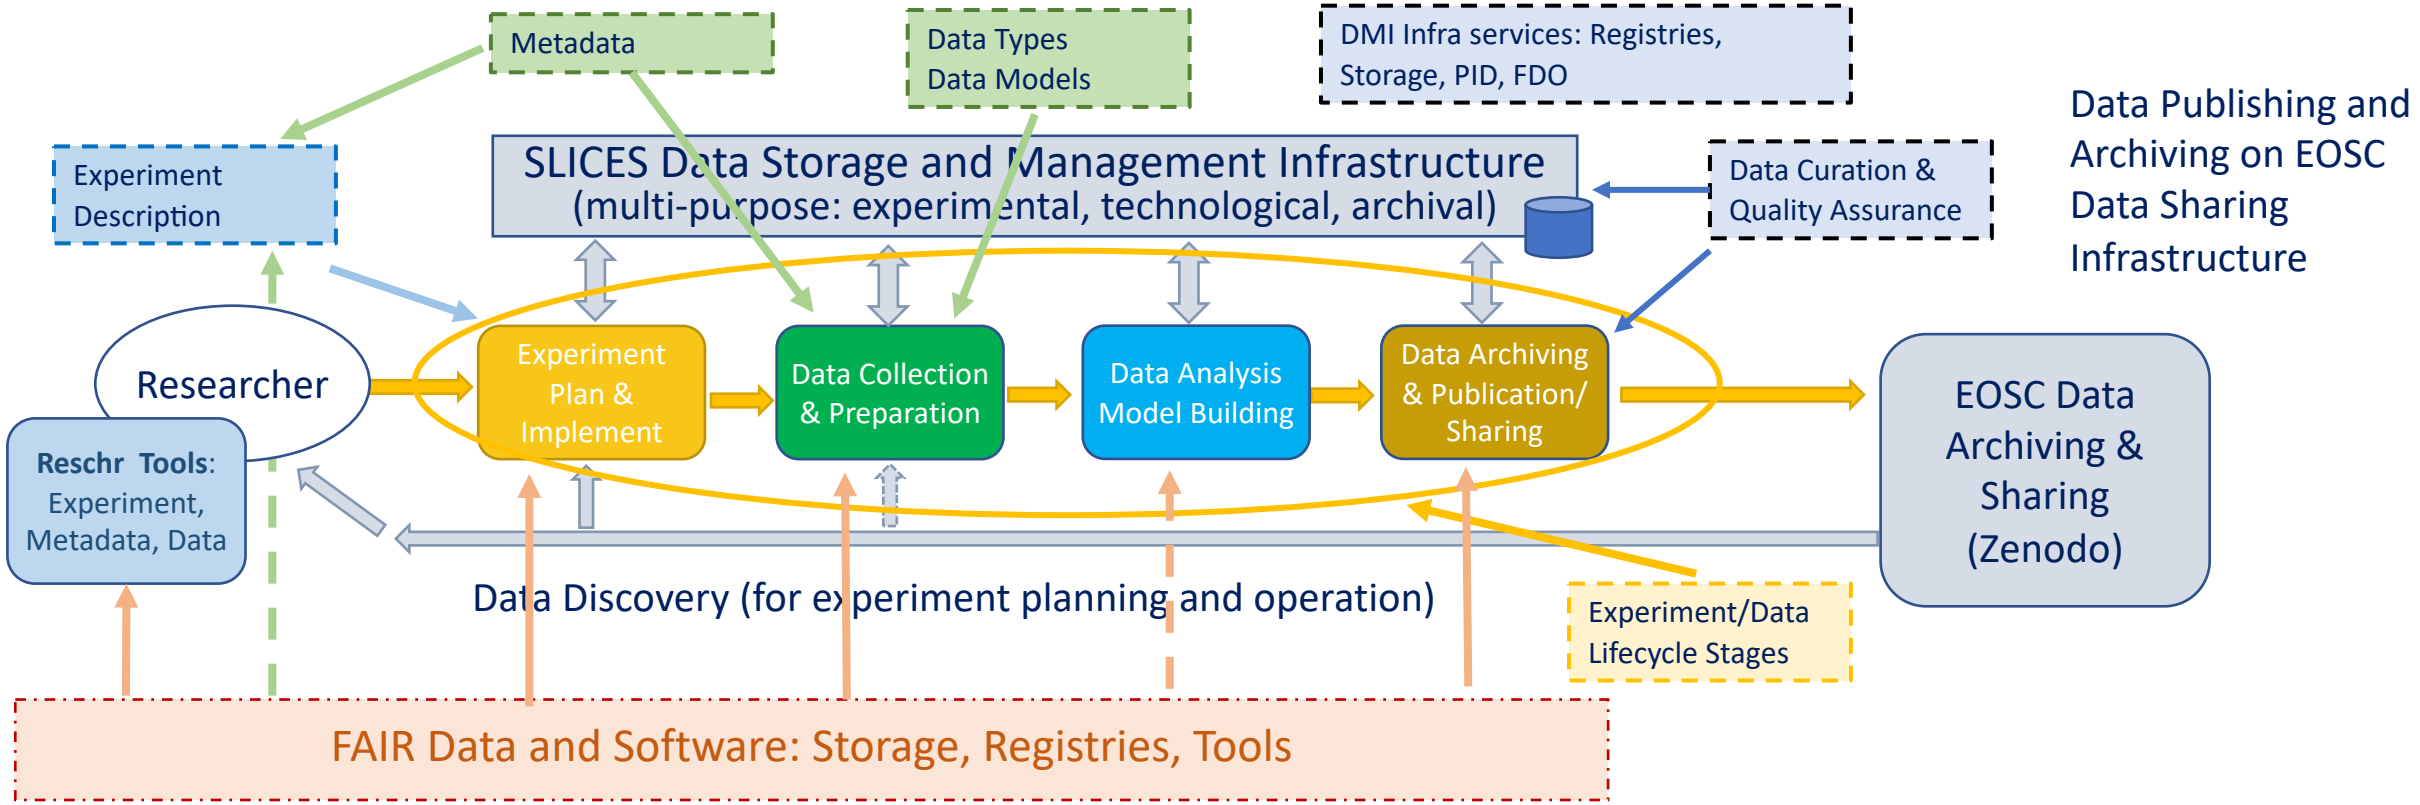

- **Each Data Lifecycle stage** – experiment, data collection, data analysis, and finally data archiving, works with own **data set,** which must be **linked**.
  - All data sets need to be stored and possibly re-used in later processes.
- Many experiments and research require already existing datasets that will be available in SLICES data repositories or can be obtained/discovered in EOSC data repositories

slicesRI

SLICES Academy

SLICES Academy

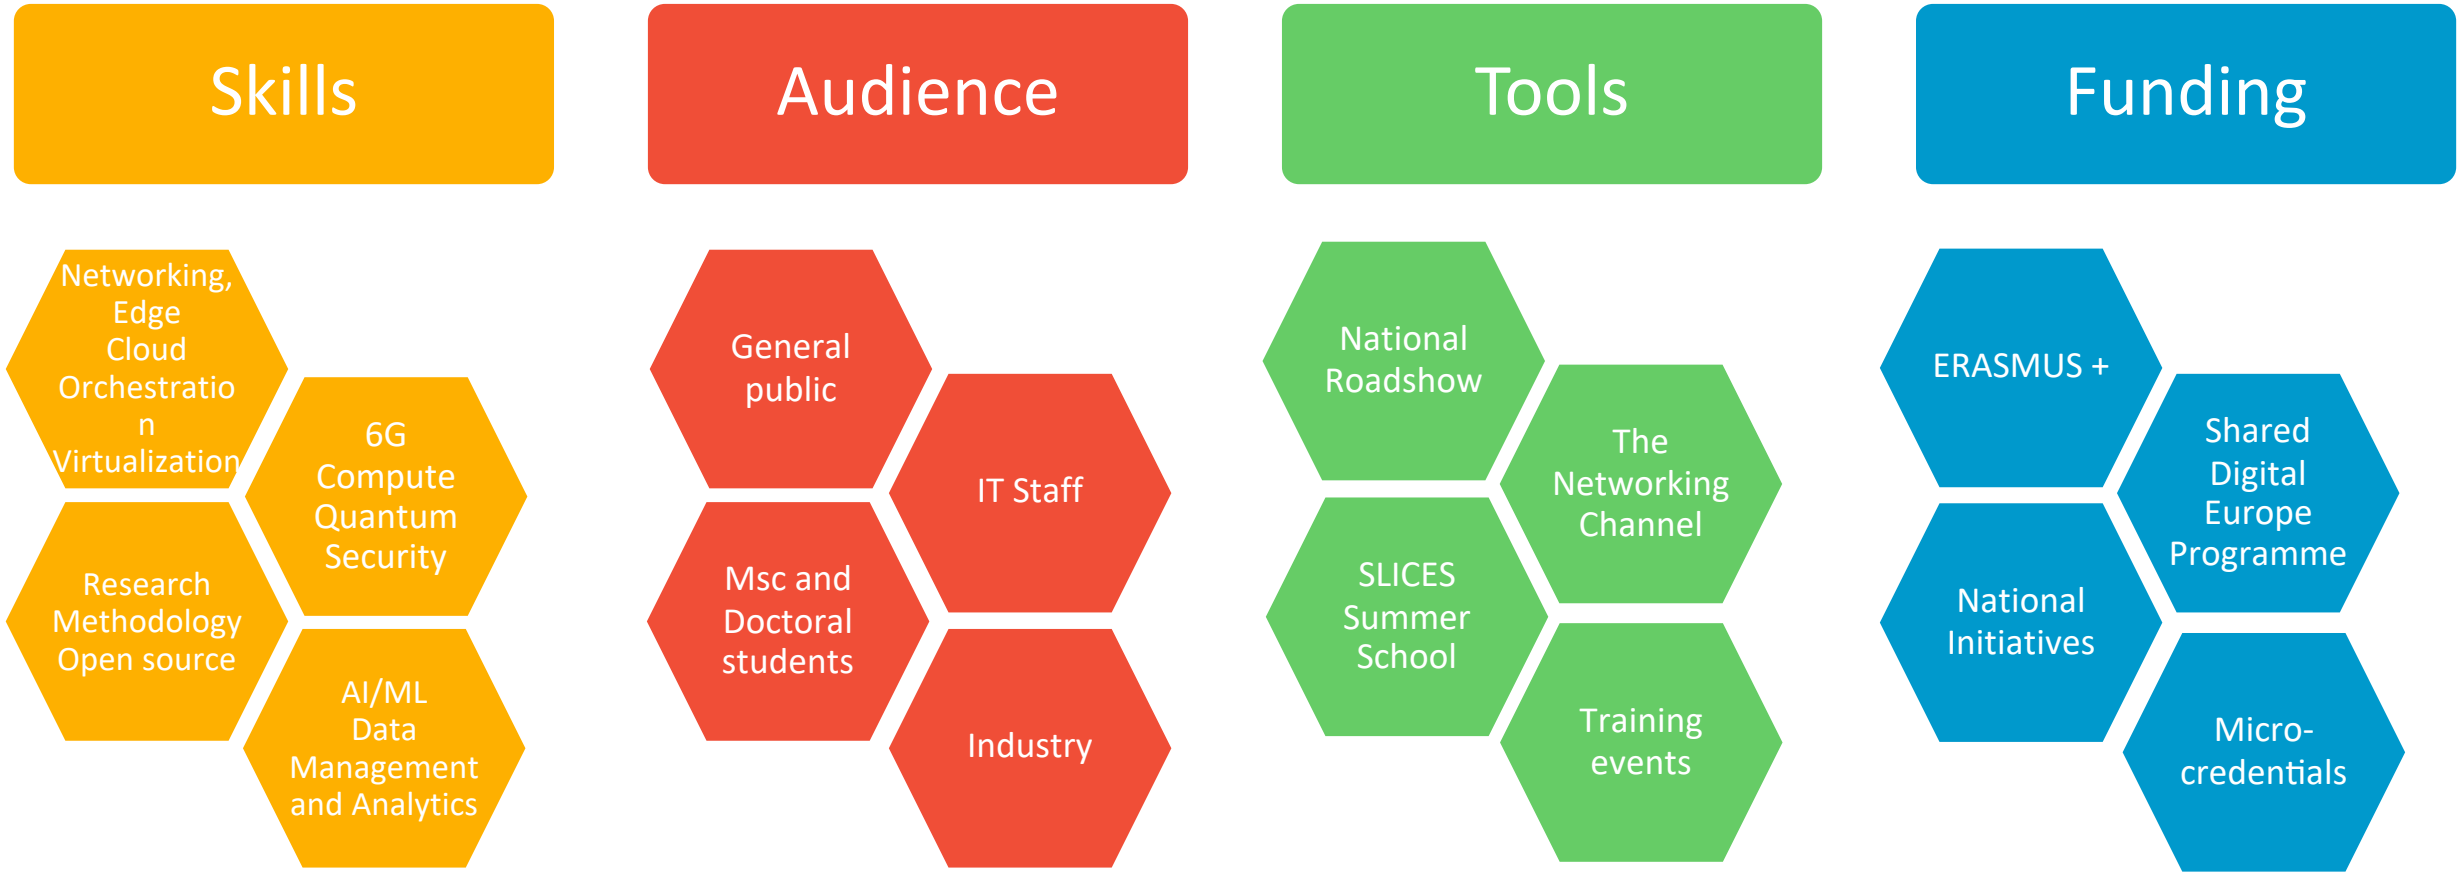

slicesRI

On behalf of SLICES consortium

Thank you

www.slices-ri.eu

For more information, please contact:
Serge Fdida
serge.fdida@sorbonne-université.fr

slicesRI