

D5.1. Intermediate report on user needs, SLICES services catalogue, access policies and training strategy

Acronym	SLICES-PP
Project Title	Scientific Large-scale Infrastructure for Computing/Communication Experimental Studies – Preparatory Phase
Grand Agreement	101079774
Project Duration	40 Months (01/09/2022 – 31/12/2025)
Due Date	30 th April 2024 (M20)
Submission Date	26 th June 2024 (M22)
Authors	Christian Perez (Inria), Andrea Passarella (CNR), Marco Conti (CNR), Eleonora Borgia (CNR), Mattia Campana (CNR), Abraham Gebrehiwot (CNR), Filippo Lauria (CNR), Andrea De Vita (CNR), Raymond Knopp (EURECOM), Paolo Bellavista (UNIBO), Armir Bujari (UNIBO), Bartosz Belter (PSNC), Sebastian Gallenmüller (TUM), Georg Carle (TUM), Brecht Vermeulen (IMEC), Eduardo Jacob (EHU), Damien Saucez (Inria), Panayiotis Andreou (UCLAN)
Reviewers	Panayiotis Andreou (UCLAN), Shashank Shrestha (UvA), Yuri Demchenko (UvA), Solenne Fortun (Inria)
Dissemination level	Public



This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101079774. The information, documentation and figures available in this deliverable, is written by the SLICES-PP project consortium and does not necessarily reflect the views of the European Commission. The European Commission is not responsible for any use that may be made of the information contained herein.



Executive Summary

This deliverable is an intermediate report of the workpackage on user needs, services, access and training strategy. It presents the current view that will be further consolidated in Deliverable D5.2 “Final report on user needs, SLICES service catalog, access policies and training strategy” (M40, December 2025) and that will provide the final view on these topics.

This deliverable starts by dealing with user needs by describing a methodology for identification of them. By applying it, three first blueprints have been identified and are presented: the “post-5G” blueprint, the “cloud/edge” blueprint, and the “machine learning/federated learning” blueprint.

The second section of the deliverable recalls the three access types (trans-national/physical access, trans-national virtual access/remote access, and virtual access) and the three access modes (excellence-driven, market-driven, and wide access) that will be supported by SLICES.

The third section of the deliverable deals with the current vision of SLICES services catalogue. Services have been further divided into supporting and basic services. As a consequence of blueprints, a new category of services, SLICES Blueprint services, have been introduced. It aims at gathering the services provided by the blueprints, i.e., services specific to a particular research community. It is important to notice that this part also considered services from a pre-operation point of view and therefore also contains implementation considerations. This work has been carried out in collaboration with relevant workpackages: WP3 “Scientific and technical strategy and specifications”, WP6 “Operational framework”, and WP7 “Data management and ethics requirements”.

The fourth section focus on training activities in particular with respect to four objectives: i) to identify the training needs and training methodologies that will be followed; ii) to develop and provide the respective training material to organize SLICES-RI training events (training sessions, webinars, plugfests, hackathons) as well as shared teaching material used in SLICES-RI for teaching basic skills at the convergence of computing and networking (SLICES Academy); iii) to provide guidelines, inter-site collaboration incentives and alignment with national programs pertaining to the teaching of key technical skills in the areas of interest of SLICES-RI (SLICES Academy); and iv) to facilitate researcher mobility, among the SLICES-RI member institutions and for the research community at large, for the exchange of know-how among the users of the facilities.



Table of contents

EXECUTIVE SUMMARY	2
TABLE OF CONTENTS	3
TABLE OF FIGURES	4
TABLE OF TABLES	4
1. USER NEEDS MONITORING AND REFINEMENT	5
1.1. METHODOLOGY FOR IDENTIFICATION OF USER NEEDS	5
1.2. EXAMPLES OF APPLICATION OF THIS METHODOLOGY.....	7
1.2.1. <i>The post-5G Blueprint</i>	7
1.2.2. <i>The Cloud/Edge Blueprint</i>	8
1.2.3. <i>Machine learning/federated learning blueprint</i>	10
1.3. THE WAY AHEAD.....	10
2. SLICES ACCESS STRATEGIES	11
2.1. ACCESS TYPES.....	11
2.2. ACCESS MODES.....	12
3. CATALOGUE OF SLICES SERVICES	12
3.1. OVERVIEW OF SERVICE TYPES	13
3.2. SLICES PRE-OPERATION SERVICES AND SLICES SERVICES.....	14
3.3. SUPPORTING SERVICES.....	14
3.3.1. <i>User and group management</i>	14
3.3.2. <i>Documentation and Online Experiment Helpdesk</i>	16
3.3.3. <i>Accountability & billing</i>	17
3.4. SLICES BASIC SERVICES	18
3.4.1. <i>Discovery</i>	18
3.4.2. <i>Resource reservation</i>	19
3.4.3. <i>Basic infrastructure Service</i>	19
3.4.4. <i>Resource configuration</i>	20
3.4.5. <i>Resource monitoring and profiling</i>	21
3.4.6. <i>Data Management Service</i>	22
3.4.6.1. <i>Storage Service</i>	22
3.4.6.2. <i>Data Management Service</i>	23
3.4.7. <i>Experiment control and orchestration</i>	26
3.5. SLICES BLUEPRINT SERVICES	27
3.5.1. <i>Post-5G Blueprint</i>	27
3.5.2. <i>Edge/Cloud Blueprint</i>	28
3.5.3. <i>Federated Learning Blueprint</i>	29
3.6. NEXT UPDATE CYCLE	30
3.6.1. <i>The QKD/PQC technology blueprint</i>	31
4. TRAINING ACTIVITIES FOR USERS	31
4.1. SLICES ACADEMY.....	32
4.1.1. MAIN OBJECTIVES OF THE SLICES ACADEMY	32
4.1.2. LICENSING OF PEDAGOGICAL CONTENT.....	33
4.1.3. ALIGNMENT WITH NATIONAL INITIATIVES	35
4.2. ORGANIZATION OF EVENTS	35
4.3. MOBILITY OF RESEARCHERS.....	37
BIBLIOGRAPHY	38



Table of Figures

Figure 1. General SLICES methodology for scientific priorities and user needs.....	5
Figure 2. Blueprint methodology and user needs.	6
Figure 3. Example of a source of a documentation webpage.	17
Figure 4. Example of member/project logs used in SLICES-SC for accountability.....	18
Figure 5. Example of experiment logs used in SLICES-SC for accountability.	18
Figure 6. Resource lease extension principle.	19
Figure 7. CLI proposal for handling services.	20
Figure 8. Example of a profiling dashboard provided by Grafana and the Kwollect infrastructure.	21
Figure 9. Data Management Infrastructure Architecture.....	24
Figure 10. SLICES Hierarchical FAIR Metadata Model (SFDO).	25
Figure 11. Federated learning service rely on multiple underlying SLICES services.....	30

Table of Tables

Table 1. SLICES Supporting Services	13
Table 2. SLICES Basic Services	13

1. User needs monitoring and refinement

1.1. Methodology for identification of user needs

The identification of the user needs is a task that needs to be carried out in synergy with the identification of the research priorities (T3.3), as the two topics are tightly intertwined together. Specifically, Figure 1 depicts the related general methodology, which has been defined in the framework of SLICES-DS project, and is reported here for the reader's convenience.



Figure 1. General SLICES methodology for scientific priorities and user needs.

Point 1 in the methodology has been the pre-requisite for the identification of the scientific dimension of SLICES overall, and has been completed during the SLICES design stage. DI here stands for “Digital Infrastructure”, which is the overall framework targeted by SLICES. Digital Infrastructures are Infrastructures built based on advanced results in the areas of Future Internet and Telecommunication research, which have the ambition of supporting transformative effects in the society. SLICES aims to be the reference Research Infrastructure for the research communities building the Digital Infrastructures of the next century. The other topics follow a cyclical approach, involving consultation with the stakeholder (2), elaboration of a research vision (3), identification of the key research areas (4) and the related priority topics (5).

The last point of the methodology (prioritised topics) is the starting point of the methodology for the identification of the user needs, which is depicted in Figure 2 in the next page.

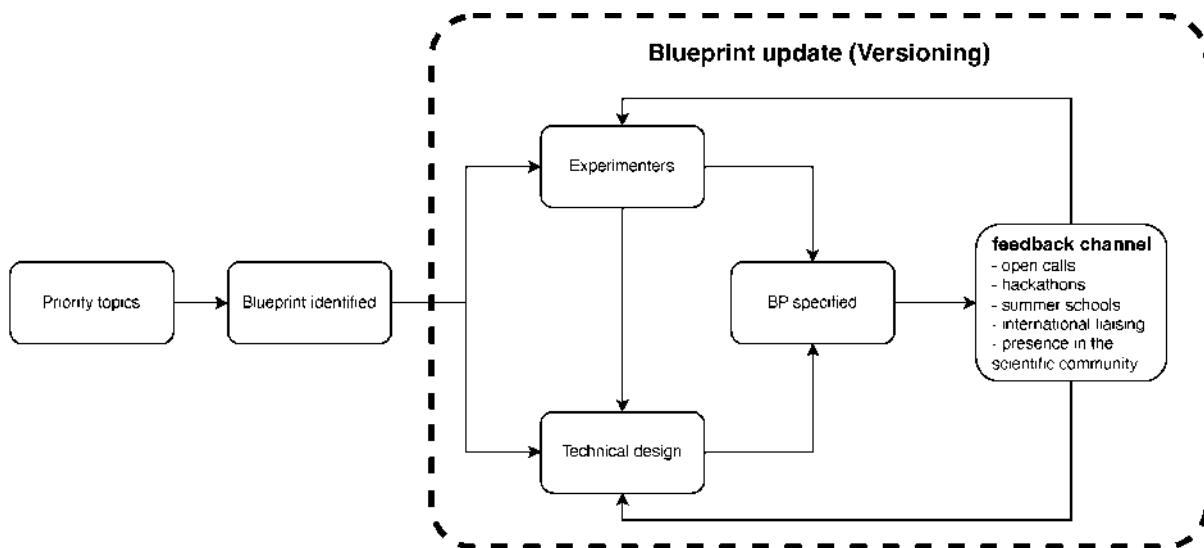


Figure 2. Blueprint methodology and user needs.

Specifically, this methodology applies to each specific research topic we identify within the SLICES framework. Each such topic is identified taking the perspective of foreseen experimenters. Specifically, we define topics considering the main priorities in the research agendas of the users' communities that SLICES members (as researchers) are part of. Typically, this is encapsulated in the concept of a dedicated "blueprint". A blueprint serves two purposes. On one hand, it encapsulates a set of services in the framework of a specific research topic. In other words, a blueprint is focused on the specific needs of a research community, and identifies the specific services that such community would need in order to run sensible experiments in SLICES. Example of blueprints are the post-5G blueprint and the cloud-edge blueprint, briefly described in Section 1.2 of this deliverable. Note that such components may be part of multiple blueprints: for example, an authentication, storage or data management component are part of many (if not all) blueprints SLICES will identify. On the other hand, a blueprint identifies the technical elements (both hardware and software) that SLICES nodes need to put in operation to support experimenters. Section 1.2 provides three examples of blueprints defined in this way.

These steps in the methodology involve both SLICES partners and users in a dynamic way. In the first place, SLICES researchers identify, thanks to their membership in the research communities they are part of, the main priorities and needs of a target community. Once such a priority is identified, an internal working group is created, tasked with the initial definition of the blueprint. This step generates a first version of the blueprint, specifying the key services that will be provided to the specific community. Then, the working groups starts deploying an initial version of the blueprint in order to make it operational.

Once this is done, the blueprint is published and advertised in the related research community. This is done using multiple channels. These channels, identified in the "feedback channel" box of the figure, implement the arrows towards both experimenters and technical designers in Figure 2. On one hand, the SLICES working group organises public events where the blueprint is presented and advertised. Summer schools focused on a specific blueprint



serve the purpose of engaging young researchers in the community as early adopters. Open calls for experiments (such as those organized in SLICES-SC) also run to open the blueprint to research groups worldwide. Last, but not least, SLICES constantly liaises with international initiatives, to discuss synergies and complementary approaches.

This phase serves the purpose of testing the functionalities of the blueprint with real users external to the SLICES consortium, in order to understand whether the design and implementation of the blueprint addresses the main needs of the target research community. Feedback is exploited to revise the blueprint technical design, and update the blueprint according to a normal versioning approach. In this way, the blueprint becomes a stable service of the research infrastructure, and evolves according to the dynamic needs of the specific research community it is addressed to.

1.2. Examples of application of this methodology

In the starting phase of SLICES (i.e., before full operation) we already identified three blueprints that are being developed according to the methodology described in Section 1.1, i.e., the “post-5G”, the “cloud/edge”, and the “machine learning/federated learning” blueprints.

The post-5G blueprint has been identified through heavy interaction in the research community working on post-5G networks (including the SNS JU). Its objective is to support experimentally-driven research in the post-5G scenario. Specifically, now that 5G is an industrial reality, the research community is looking at an evolution primarily characterised by disaggregation, modularity, and flexibility, in addition to the incorporation of novel radio technologies that will be part of future generations (towards 6G).

The cloud/edge blueprint has been identified by interacting with the community working on virtualisation of resources and resource management in the cloud-to-edge continuum. It addresses scientific evolutions from well-established cloud technologies, towards complete decentralisation of resource provisioning at the edge and beyond, possibly also considering the constraints of resource-constrained edge devices such as Internet-of-Things gateways.

In the following section, we summarise the main status of each of them, focusing specifically on the key functionalities addressing the identified user needs.

1.2.1. The post-5G Blueprint

The post-5G blueprint has been the first one identified and developed in SLICES, and it has been used to better define the various steps of the described methodology.

The full description of the blueprint is available in a dedicated part of the SLICES portal¹.

With respect to the user needs, the first version of the blueprint addresses the following ones:

¹ https://doc.slices-sc.eu/blueprint/5g_blueprint.html, Last accessed: 2024-06-06.



- *Flexibility of development options.* The post-5G blueprint allows for deployment in a range of configurations, from the simplest one, where all 5G components are virtualised and run in a single cluster, to the more realistic (and complex) one, where different components run either on bare metal or in virtualised components across multiple clusters. This allows experimenters to trade between complexity and abstraction (and simplicity of deployment).
- *Multiple radio units.* The blueprint supports multiple radio technologies, allowing for experimentation across them or within each of them in isolation.
- *Core/RAN separation.* The blueprint allows to instantiate core and RAN components of a 5G network in multiple ways, also supporting virtualisation of core functions. The various logical modules of each segment can be deployed according to appropriate virtualisation technologies.
- *Connectivity.* The blueprint supports multiple network configurations across the various logical modules of the Core and RAN. It supports both scenarios where dedicated connectivity is available (e.g., inside a single operational infrastructure) or where components are connected through research networks (typically, and NREN) or even via the public Internet.

We are now in the process of defining the third version of the 5G blueprint, after almost two years of testing and interaction with the research community based on the first delivered version. Some of the main features foreseen for V3 are as follows:

- *RIC support and full O-RAN compliance.* The blueprint is designed having in mind the O-RAN framework, whereby all functions are modular, and can be provided by different entities according to well-defined APIs. In addition, a specific focus for V3 of the blueprint is full support for the RIC functions in the O-RAN model.
- *Data management.* SLICES is developing a complete solution for experimental data management, compliant with EOSC² requirements, such that SLICES data can be exported in EOSC-compliant systems. V3 of the post-5G blueprint will provide full support for this function.
- *Reproducibility of experiments.* The blueprint is gradually incorporating a framework for reproducible experiments (Plain Orchestrating Service (POS)), which is customisable for the specific need of 5G experimentation.

1.2.2. The Cloud/Edge Blueprint

The Cloud/Edge blueprint has been identified in the second part of 2023 as the second blueprint to be developed in SLICES, in order to target the needs of the research communities working on virtualised networking/storage/computing resources on multiple heterogeneous clouds, on hybrid public/private clouds, on intermediary nodes, on Multi-access Edge Computing (MEC) nodes, and on on-premise gateways, e.g., for IoT sensors/actuators integration in Cyber-Physical Systems experimentation. In short, in the following we will use the term cloud-to-edge continuum to identify this set of variegated deployment scenarios.

² European Open Science Cloud (EOSC) is developed and positioned as European Federated Data Sharing infrastructure for European research. EOSC provides a set of services and tools to share and publish research data.



A dedicated SLICES discussion group has worked on the definition and on a preliminary description of the Cloud/Edge blueprint, which is expected to be finalized (in its V1 version) and made available in the SLICES portal by the end of June 2024. In addition, the Cloud/Edge blueprint V1 will be disclosed to the research community and opened for experimentation/feedback (according to the methodology depicted in Figure 2) during the SLICES summer school in July 2024.

The first version of the Cloud/Edge blueprint will address the following user needs:

- **Flexibility of development options.** The blueprint will support a range of different development/deployment configurations, as it is crucial for the variety of scenarios of interest for the cloud-to-edge continuum community. It will consider simple scenarios with a single remote cloud-based cluster of virtualised resources hosting all the experiment components except for data fusion components deployed locally to sensors/actuators, but also more articulated ones where multiple cloud providers (public and private) host more computing/storage-intensive tasks, while others are more distributed over multiple Multi-access Edge Computing (MEC) nodes and edge gateways, possibly by supporting the dynamic migration of experiment components in the cloud-to-edge continuum. As for the post-5G blueprint, this allows experimenters to trade between complexity and abstraction (and simplicity of deployment).
- **Support for multiple cloud providers.** The blueprint will consider and support multiple public/private cloud providers, either on-premise or geographically distant, allowing for experimentation with virtualised resources across them or within each of them in isolation. In addition, the blueprint will enable the cloud-based hosting not only of 5G core and RAN components, but also of application-level components that are integrated in the experiments (e.g., for computing-intensive training of machine learning models). Moreover, heterogeneous cloud support stacks will be integrated as possible options in the blueprint to offer experimenters the opportunity of differentiated abstraction levels to access the available bare-metal resources.
- **Inter-cluster orchestration.** The blueprint will allow to instantiate experiment components over remote clusters of virtualised resources on different cloud providers and intermediary nodes in the cloud-to-edge continuum. Inter-cluster orchestration will be developed as the federation and extension of existing state-of-the-art orchestrators in order to maximize code re-usability and to minimize the experimenters' efforts towards new tools.
- **Hierarchical resource composition.** The targeted Cloud/Edge research communities recognize the need for managing and coordinating experiment resources at different layers, from application/service deployment to infrastructure provisioning. The blueprint will support and simplify this via descriptions of resource compositions, called bundles, for tiered (layered) deployment, at different levels of abstraction, also made available as templates/examples in a centralized repository. This support will consider the integration with declarative control plane frameworks such as Crossplane³.

³ Crossplane.io, <https://www.crossplane.io/>



- **Data management.** The blueprint considers supporting experimentation for machine learning (in particular, federated learning) on IoT data as a central need for the targeted communities of experimenters. Therefore, in synergy with what is done for the post-5G blueprint, the Cloud/Edge blueprint will integrate with the SLICES framework for experimental data management and will be EOSC-compliant.
- **Reproducibility of experiments.** In addition, this blueprint is gradually incorporating the same framework for reproducible experiments developed by TUM, customised for the specific multi-deployment needs of Cloud/Edge experimentation. This functionality will be made available in the V2 of the blueprint.

1.2.3. Machine learning/federated learning blueprint

Machine learning/federated learning has been used, starting 25/05/2023, as one of the possible services in the Task 3.1 “architecture long term vision” meetings. Since the virtual plenary in July 2023 it has also been on the agenda for pre-operations.

During the discussions in Task 3.1 “architecture long term vision” and later on in specific WGs as basic infrastructure and storage, we identified more use cases and user communities (for now internal with partners of SLICES, but not the people creating the infrastructure). We decided to start from Sept/Oct 2024, after the start of pre-operations with the first version of the federated learning service and its accompanying services such as basic infrastructure and storage, with a regular dedicated brainstorm meeting how to broaden up this blueprint to a larger community also outside of SLICES partners.

The basic goal of this blueprint is to identify and define the services needed for advanced and easy experimental (federated) machine learning research. For the first version, this has led to the definition of the following services:

- Basic infrastructure service: to be able to use bare metal servers, virtual machines and embedded/IoT devices;
- Widely accessible S3 storage service;
- The federated learning service itself (using the other mentioned services).

And with input to the following services:

- Interconnectivity service: service to interconnect infrastructures on layer 2;
- GPU Lab infrastructure (GPU Lab/jupyter notebook⁴) for easy use of GPUs with machine learning.

For next versions, it might be interesting to also integrate reproducibility and data management services in the same framework as the two other blueprints, although the foundations for those are already embedded and being discussed.

1.3. The way ahead

⁴ <https://doc.ilabt.imec.be/>, Last accessed: 2024-06-06.



The methodology described in this deliverable is going to be used to support other blueprints in the pre-operation phase. This will allow us to enlarge the user base of the research infrastructure, and obtain additional feedback from the relevant research communities with respect to the overall approach of service provisioning in SLICES.

Specifically, in addition to evolving the three blueprints described in this deliverable, we anticipate to launch specific working groups for the design of blueprints in the following areas:

- Integrated sensing and communication;
- Pervasive Artificial Intelligence;
- Quantum Networking and Computing;
- IoT;
- Security;
- Sustainability and Energy efficiency.

More details on the specific features identified so far in each of these domains, according to the described methodology, are presented in D7.1 “Report on Future Evolution of the RI” of the SLICES-SC project.

2. SLICES Access strategies

The access types and modes defined in SLICES DS D2.2 [DS-D2.2] are the starting point for SLICES-PP discussion. Up to now, we do not have identify any needs to update them. There are recalled in the remainder of this section.

2.1. Access types

Three types of access will be enabled through SLICES, in compliance with the access methods defined in the ESFRI 2020 white paper for “MAKING SCIENCE HAPPEN - A new ambition for Research Infrastructures in the European Research Area” [ESFRI20] as follows:

- **Transnational access (Physical access):** Apart from accessing the equipment, the facilities and the laboratories, users will also be offered with the required technical and scientific assistance to learn and use the infrastructure. Though such a type of access is more applicable to RIs that cannot be easily accessed remotely for fine-tuning, this is envisioned in order to facilitate experiments requiring physical presence at one of the SLICES sites, such as for instance Bring-Your-Own-Equipment (BYOE). In such experiments, researchers can install their prototype equipment at one of the SLICES sites and combine it in their experiments with the equipment offering of SLICES.
- **Transnational virtual access (Remote access):** The majority of the tools and services that are designed and shall be developed for SLICES deal with the remote access method, in order to present to the users of the platform a unified solution for retrieving, selecting, reserving their experimental components and deploying their experiment on top of them. Through remote access, the users of the platform will take advantage of the tools for controlling their experiment and the environment



parameters in an organized manner, under real world settings. This type of access is envisioned as the preferred method of access for the vast majority of SLICES users.

- **Virtual access:** Virtual access typically concerns access to data and digital tools. The provision of Virtual access to SLICES is aided through sophisticated cloud services and communication networks and allows for the remote access to repositories and archives of produced experimentation results. The vision of SLICES for allowing virtual access is to offer a pan- European operational networking and computer infrastructure to facilitate scientific research with instrumentation and experimentation capabilities.

2.2. Access modes

As presented in the previous section, users will mostly access SLICES-RI in a remote fashion (remote and virtual access modes) via Internet. Hence, most SLICES services are computer provided services accessible through a dedicated web portal, dedicated APIs (for example through a REST software architecture), and, also in order to support experimentation by users from research domains different from future Internet and distributed systems, high levels tools such as Jupyter notebooks and workflow systems. Based on the expertise of the SLICES partners in operating research infrastructures, SLICES will provide the three access modes described in the document “European Charter for Access to Research Infrastructure” [EC16] of the European Commission:

- **Excellence-driven** access: the access is exclusively dependent on the scientific excellence, originality, quality and technical and ethical feasibility of an application evaluated through peer review conducted by an Access Committee.
- **Market-driven** access: the access is based on a fee and an agreement between the User and the RI.
- **Wide access:** the access is available to any eligible users. It may be subject to stronger quotas/limitation than the two other access modes.

Following the analysis of Deliverable D1.1 “Technological status and capabilities of existing ICT Research Infrastructures” [DS-D1.1], large academic experiments are expected to access the platform under the “excellence-driven” mode while smaller scale experiments will benefit from the “wide” mode. Business and industry are mainly expected to access the platform under the “market-driven” mode. Considering the user community, an initial estimation is that “excellence-driven” mode will represent 60% of the time platform, “wide” 20%, and “market-driven” 20%. These numbers will be further refined and managed by the CMO based on the recommendations of an Access Committee in relation with the other Committees and under the decision of the Supervisory Board. Similarly, the prices to access the platform for the “market-driven” mode are yet to be discussed and decided. The two other modes, excellence-driven” and “wide”, are open and free of charge for European academics. If needed, calls for proposals—such as those organized by PRACE for example— will be organized to allocate resources based on best usages by the Access Committee.

3. Catalogue of SLICES Services



3.1. Overview of Service Types

Deliverable D2.2 of SLICES-DS “SLICES as a Service, baseline” [DS-D2.2] found twelve services that were grouped into four categories of services from a user point of view. In SLICES-PP, these services have been further grouped into two categories: Supporting Services and Basic Services. With respect to SLICES-DS D2.2, Supporting Services (Table 1) gather services that belong to the User and platform management services while Basic Services (Table 2) covers Resource management services, Data oriented services, and Experiment management services.

Table 1. SLICES Supporting Services

	Type of Service	Working Group
User and platform management services		
USERS_MGT	User and group management	User management service
DOCUMENTATION	Documentation and Online Experiment Helpdesk	User management service
ACCOUNT	Accountability & billing	User management service

Table 2. SLICES Basic Services

	Type of Service	Working Group
Resource management services		
DISCOVERY	Resource discovery and description	Data Management Infrastructure
RESERVATION	Resource reservation	Basic Infrastructure service
CONFIGURATON	Resource configuration	Basic Infrastructure service Connectivity service
MONITORING	Resource monitoring and profiling	Monitoring Service
Data oriented services		
DATA	Data Management Service	Storage service Artefact deposit service Data Management Infrastructure
ANALYSIS	Experiment data validation and correlation with other experiments	-
Experiment management services		



EXP_MGT	Experiment management	-
ORCHESTRATION	Experiment control and orchestration	Plain Orchestration Service
DASHBOARD	Dashboard	-

In this deliverable of SLICES-PP, we have considered in more details nine services of the 12 service types identified in SLICES-DS. This section aims at gathering the current understanding of the SLICES-PP service catalogue, as well the level of definition of these services. The updated view come from various SLICES-PP working group ranging from service level working groups to blueprint working groups but also architecture (WP3) and data (WP7) workpackages.

A major evolution of SLICES service catalogue is that community-oriented services (that we call “Blueprint services”, following the methodology described in Section 1) have been identified as a new category of services in SLICES-PP. This category contains currently the blueprints services introduced in Section 1.2 and described in more details hereafter the Supporting and Basic Services.

3.2. SLICES Pre-operation Services and SLICES Services

In relation with other workpackages, in particular with WP3, WP4 and especially WP6, services to be considered and included into the pre-operation platforms have received more attention, in particular from an implementation point of view. That is why this deliverable not only describes the services from a user point of view but also mentions information from an implementation point of view when possible.

3.3. Supporting Services

This section deals with supporting services that contains 3 types of services as described in Table 1

3.3.1. User and group management

For determining the vision about user and group management and a portal to be used by the experimenters, a working group was started within the WP3 architecture group. We started with listing up requirements for this, split in a pre-operations list and a list for operations, and subdivided into technical and non-technical requirements. These are based on the experience of user management in Fed4FIRE (10 years) and SLICES-SC (2 years), with new insights from the working group. These were finished in the meeting of March 1st 2024.

Requirements pre-operations:

- Non-technical:
 - There is no official ERIC entity yet, so the partners participating in the pre-operation phase should take the responsibility, e.g. for privacy policies, etc.
 - Lessons learned:



- Define the info fields that you want to know from users (name, email, student/academic/industry, company or institution, city, country) and from projects (project name, what do you want to do, which infrastructure do you want to use, how did you hear about us).
- Need for approval procedure (approving “projects” with a PI (Principal Investigator) and delegate to him/her for further approval of users in the project works very well. Roles of admin, project lead, project admin and user.
- Separation between user and project works very well. Expiration date on projects (you can ask for an extension), and a user chooses in which project s/he runs experiments. If a user is not in an active project, s/he cannot start an experiment.
- Edugain for academic accounts works quite well (including 2FA), some universities do not expose the needed fields (e.g. email), but it could be adapted when asked.
- A lot of automatic emails need to be supported (approval, extension, PI delegation, password reset...).
- Management always asks statistics about accounts/experiments.
- Technical:
 - OAuth2 and OpenID Connect provider. It allows Clients (also known as Relying Parties) to verify the identity of the End-User, as well as to obtain basic profile information about the End-User in a REST-like manner.
 - Access_token and id_token returned by the OIDC login flow are JWT tokens containing information on the end user.
 - Documentation is available online⁵.
 - An open question is whether we need an integrated services dashboard for pre-operations.

The **requirements for operations** are a bit more vaguely defined as several options need to be investigated.

Non-technical requirements/options for operations:

- There will be an official entity (ERIC) that operates and manages users and is responsible.
- European charter of access for research infrastructures [EC16] provides many information:
 - Workflow probably to evaluate through peer review, see also access modes (excellence-driven access, market-driven access, wide access).
- An option is to consider Geant eduTEAMS⁶.
- Shall we allow two level of users: node level users and SLICES level users? If yes, we need to investigate whether is it legal and wanted to export users from a local site to the central hub.

⁵ https://doc.slices-sc.eu/testbed_owner/oauth.html, Last accessed: 2024-06-06.

⁶ <https://www.eduteams.org/> and <https://wiki.geant.org/display/eduTEAMS/Documentation>, Last accessed: 2024-06-06.



- Switzerland provides Switch edu-ID which is a universal login for lifelong learning⁷.
- Policies trans-national access, trans-national virtual access, virtual access.

Technical requirements for operations:

- Central or also delegated users from nodes:
 - API to create users?
 - Local administrator.
 - Approve users per node/country by delegation?
- Expiring users over time.
- Are social authentication methods supported (e.g., Google, Facebook)? It seems difficult because the status of academic shall then be established by an additional mechanism.
- Make sure platforms can scale to thousands of users
- Live replication in 2 sites in 2 different nodes (countries): we should separate NREN networks in case something happens with a full NREN network (has happened in several countries in the last years).
- Can the portal track the audit trails and logged user activities? It may be required for auditing and compliance purposes and also for user behaviour analytics.
- We need to investigate whether we will develop the software stack from scratch or start from existing implementations.
- SLICES needs to support scheduled experiments (per service if needed).
- Policy documents need to be created.

3.3.2. Documentation and Online Experiment Helpdesk

For a documentation website, we have very good experience with a back-end in Git/GitLab where multiple people can edit easily, and automatic publishing to a frontend, e.g. with Sphinx⁸. Advantages of this are:

- Easy editing by multiple people ((e.g., in markdown or in RST));
- Traces and history of changes with the standard git tool;
- Gitlab allows also web editing which is easy for small changes.

Example can be seen on online documentation⁹.

As a detailed example, the page¹⁰ is defined by the source in Figure 3 (screenshot of GitLab).

An online experiment helpdesk should be reachable by email and be backed by a ticketing system (which can also allow web access to the tickets). Ideally, it is also easy to submit tickets/questions from within tools/services, which allow e.g. to add track traces or more detailed information which is helpful for the helpdesk team. As such, an API should be available for the tool developers.

⁷ <https://www.switch.ch/en/edu-id>, Last accessed: 2024-06-06.

⁸ <https://www.sphinx-doc.org/>, Last accessed: 2024-06-06.

⁹ <https://doc.slices-sc.eu/>, Last accessed: 2024-06-06.

¹⁰ https://doc.slices-sc.eu/testbed_owner/oauth.html, Last accessed: 2024-06-06.



In SLICES-SC it was organized as described on SLICES-SC’s Support Page 11 where the helpdesk email was backed up by Atlassian service desk/JIRA (with API, accessible e.g. by the jFed tool).

As Large Language Models are evolving quickly, we shall investigate whether/in which cases a chatbot could be useful for SLICES users.

```
oauth.rst 3.42 KB
1 Federating via OAuth2
2 =====
3
4 The Slices portal is a standards-compliant OAuth2 <https://oauth.net/2/> and OpenID Connect provider <https://openid.net/connect/>. It allows
5
6 OpenID Connect extends OAuth2. The OAuth2.0 protocol provides API security via scoped access tokens, and OpenID Connect provides user authentication
7
8 The Slices portal provides all OpenID Connect metadata via <https://portal.slices-sc.eu/.well-known/openid-configuration>.
9
10 The :code:'access_token' and :code:'id_token' returned by the OIDC login flow are JWT tokens containing information on the end user. The JWKS contains
11
12 Registering your service as an OAuth Client
13 =====
14
15 Please email helpdesk@ilabt.imec.be with the following information on your service:
16
17 * Service Name
18 * Service URL
19 * One or more redirect URI's: valid callback locations where the authorization code or tokens should be sent to. (**Tip:** also include any redirect
20 * Token endpoint auth method: either :code:'client_secret_basic' or :code:'client_secret_post' are supported.
21
22 We will then register your service in the portal and provide you with a :code:'client_id' and :code:'client_secret' to use.
23
24 Familiarizing yourself with the authentication workflow
25 =====
26
27 If you want to familiarize yourself with the OIDC workflow, you can use either the 'Open ID Connect Playground' <https://openidconnect.net/> or '
28
29 Note that the provided Client ID and Client Secret only have 'https://openidconnect.net/callback' and 'https://hoppscotch.io' as a registered redirect
30
31 OpenID Connect Playground
32 =====
33
34 1. Click on the 'Configuration' button
35 2. Select Server Template 'Custom'
36 3. Use Discovery Document URL <https://portal.slices-sc.eu/.well-known/openid-configuration>
37 4. Set the OIDC Client ID to :code:'zj5cm8rJBW7j9v5sUeVF3sv' and OIDC Client Secret to :code:'PEesNzJrtjXLoSLlDXVfWDFRmAIHryvB3fRhdEKebBHKes2V'
38 5. Set Scope to 'openid userinfo'
39 6. Save the configuration
40
41 You can now perform the full authentication flow:
42
43 - Redirect to OpenID Connect Server
44 - Exchange Code from Token
45 - Verify User Token
46
47 Hoppscotch
48 =====
```

Figure 3. Example of a source of a documentation webpage.

3.3.3. Accountability & billing

This has currently not been discussed in depth, as it is more depending on policies for operations. What is currently being used in SLICES-SC for accountability, can be seen in Figure 4. Example of member/project logs used in SLICES-SC for accountability. and Figure 5. Example of experiment logs used in SLICES-SC for accountability., where we save tracks of user and project roles and the experiments created by users.

¹¹ <https://doc.slices-sc.eu/support.html>, Last accessed: 2024-06-06.



2023-06-13 13:12 CEST	Created project summerschool2 with lead Brecht Vermeulen
2023-06-13 16:38 CEST	Brecht2 Vermeulen Added member Brecht Vermeulen in role MEMBER to slice dem1
2024-02-20 11:52 CET	Mathias De Brouwer Added member Brecht Vermeulen in role MEMBER to project dp2024-11
2024-04-23 02:25 CEST	Created project test_onderstreek with lead Brecht Vermeulen
2024-04-23 02:25 CEST	Updated member Brecht Vermeulen attribute portal_home from slices to imec
2020-06-29 09:07 CEST	Thijs Walcarius Added member Brecht Vermeulen in role MEMBER to slice enos1
2020-09-28 20:59 CEST	Created project icn2020 with lead Brecht Vermeulen
2020-09-29 18:11 CEST	Thijs Walcarius Added member Brecht Vermeulen in role MEMBER to slice enos1
2023-04-07 10:21 CEST	Changed member Brecht Vermeulen to role ADMIN in project slices-tut
2024-03-06 02:03 CET	Created project slices5gcore with lead Brecht Vermeulen
2021-09-09 16:08 CEST	Removed member Brecht Vermeulen from slice strat3
2021-09-09 16:08 CEST	Removed member Brecht Vermeulen from project stratum2
2021-09-09 16:08 CEST	Removed member Brecht Vermeulen from project stratum
2021-09-09 16:08 CEST	portal Added member Brecht Vermeulen in role MEMBER to project stratum

Figure 4. Example of member/project logs used in SLICES-SC for accountability.

Slices

Slice name	Project	Owner?	Created	Expiration
testbv	thesis-liesbeth	Yes	2021-05-05 10:31 CEST	2021-05-05 12:31 CEST
testbv2	thesis-liesbeth	Yes	2021-05-05 10:36 CEST	2021-05-05 12:37 CEST
testbv3	thesis-liesbeth	Yes	2021-05-05 11:02 CEST	2021-05-05 13:02 CEST
testbv4	thesis-liesbeth	Yes	2021-05-05 11:03 CEST	2021-05-05 13:04 CEST
cont1	bvermeul	Yes	2021-09-08 18:16 CEST	2021-09-08 20:16 CEST
strat1	stratum	Yes	2021-09-09 12:50 CEST	2021-09-09 14:50 CEST
strat2	stratum	Yes	2021-09-09 12:52 CEST	2021-09-09 14:52 CEST
strat4	stratum	Yes	2021-09-09 16:09 CEST	2021-09-09 18:09 CEST
demo1	bvermeul	Yes	2023-01-26 10:03 CET	2023-01-26 12:03 CET

Figure 5. Example of experiment logs used in SLICES-SC for accountability.

3.4. SLICES Basic Services

This section deals with SLICES Basic Services that contains 3 groups of services as described in Table 2 for a total of 9 services. At this stage, only 6 services have been further analysed.

3.4.1. Discovery

The Discovery Service is in charge of determining which resources are available on each site, and reporting back to the central Hub. Through there, the experimenter shall be able to list the available resources that can be included in an experiment instance, and the site provider will be able to determine the utilization of resources, as well as to list which resources are malfunctioning for taking corrective actions. The service shall consist of two different parts, one at the central Hub location that will be in charge of polling/collecting the inputs from all the sites, and a client version (located at each site) that will provide the information to the central Hub. A common API will be developed for all the sites, allowing listing the available/non-available machines, the current occupancy of the machines, and any other



monitoring information that needs to be reported back to the central Hub (depending on the node usage and capabilities).

3.4.2. Resource reservation

Resource reservation should be possible instantly (i.e., a “future” reservation starting immediately) and in the future. Especially for larger experiments or experiments with scarce resources (e.g. international L2 bandwidth or special unique resources such as a radio spectrum) reservations in the future are needed. This depends however on the service. For example, for a basic infrastructure service offering virtual machines, it is much easier to create resources. However, if an experiment wants to use 1000 virtual machines, a reservation in the future might be needed as such amount of resources may not be immediately available. Each service will have to define when the reservation procedure is needed. In Figure 6, you can see how leases for resources can also be extended (if available). The tool mechanism and terminology should be very similar for the multiple services.

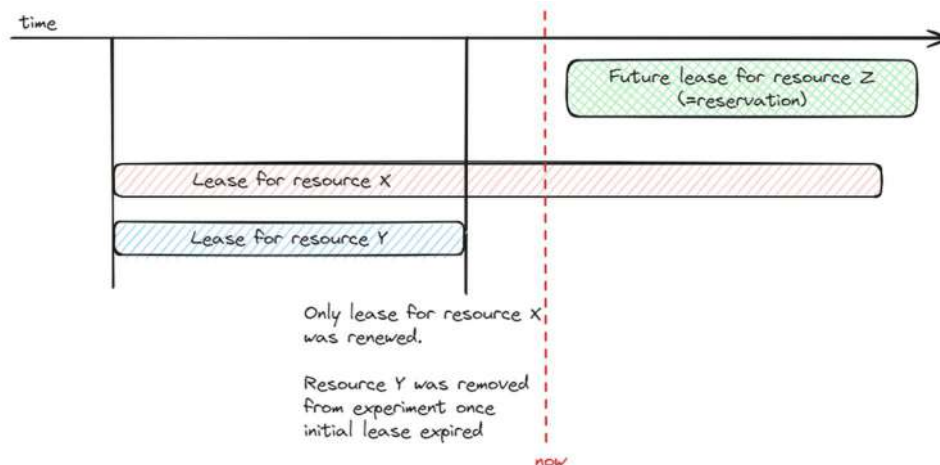


Figure 6. Resource lease extension principle.

A particular case of resources concerns the network between sites that are managed by the Connectivity Service (CS). It will expose APIs towards other SLICES services to allow reservation of network resources available for experimenters. From the point of view of the CS a resource is either a VLAN tag, which identifies the entry point to a virtual network for an experiment (realized as a GÉANT+ service over the GEANT network) or a VNI, the VXLAN network identifier for an overlay VXLAN network over the public Internet.

3.4.3. Basic infrastructure Service

The basic infrastructure service has been defined as a service with two goals:

- Giving experimenters low level access to hardware (e.g. root access on a bare metal server or virtual machine, programming firmware on an embedded device, etc.);
- Be a building block for more integrated services.

The following 2 examples should make it clear what exactly is meant:



- If an experimenter wants to use a kubernetes cluster (similar for openstack, etc.), she has two options: 1) use bare metal servers or virtual machines and deploy kubernetes. This makes that she will have full access and control and can actually change kubernetes. 2) use a SLICES kubernetes service that deploys a kubernetes cluster for the experimenter. The experimenter has no access to the below hardware. This service can use basic infrastructure services or have its own infrastructure. For 1), the experimenter is using the basic infrastructure service. For 2) she is using another service.
- An experimenter wants to use storage. If she only wants to store data from an experiment, she can use e.g. the S3 based storage offered by SLICES. If however, she is interested in analysing storage performance and comparing multiple setups, she might use basic infrastructure bare metal infrastructure with disks and set up everything herself.

The basic infrastructure service that we are currently defining, will offer low level access to bare metal servers, virtual machines and embedded devices.

3.4.4. Resource configuration

Resource configuration will of course depend on the type of resources (and hence the services). Nevertheless, we envision that we can have the same structure for all services:

- A client library for at least python;
- A CLI using that client library;
- A console web interface.

Some services might then have specific GUIs that make it easier to use the service.


We aim also to have the CLI in a basic SLICES structure. The current proposal from the basic infrastructure working group is described in Figure 7.

- One global ‘slices’ CLI, which can call other executables when a subcommand is specified:

```
$ slices bi action arg1 arg2
```

Causes the following executable to be called:

```
$ slices-bi action arg1 arg2
```



Goal 1: One CLI acts as entrypoint for manipulating all Slices resources

Goal 2: Centralize credential management

Goal 3: Keep flexibility on actual implementation, as every team can choose their own technology stack for implementing the CLI

Figure 7. CLI proposal for handling services.

The Connectivity Service (CS), under development within WP6 of SLICES-PP, in its final release will provide a set of services for experimenters to setup connections between different sites participating in a specific experiment. The CS will offer two types of services for experimenters:



- L2 connectivity over the GÉANT network utilizing GÉANT+ services realized as an LDP-signaled and ISIS-SR transported point-to-point layer 2 circuit, and,
- L2 connectivity over the public Internet, realized as an overlay VXLAN network.

From the point of view of resource configuration, the above services will rely on a pre-configured network setup, being seen as a network infrastructure ready to use by the experimenters. The CS will not configure any resources on the path between involved sites, crossing over the intermediate independent networks, i.e., campus networks, NRENs and GÉANT. However, there is a need to configure last mile networks involved in the experiment, i.e. local network infrastructures at sites.

In the first rollout of the Connectivity Service, as part of pre-operation, it is not planned to develop tools and services to configure local network infrastructures. However, the WP6 is aiming at carrying out a set of internal experiments to learn about the actual needs and requirements for local network configuration. This way lessons learned from early pre-operation experiments can be accommodated in the second rollout of the Connectivity Services in the pre-operation phase.

3.4.5. Resource monitoring and profiling

The Monitoring and Profiling service is designed to capture experimental data, collect them and to expose them through some interfaces, frameworks, and tools. This service allows research communities to analyse obtained results, reproduce, and observe some experiments. This service should also allow infrastructure managers and providers to observe deployed services and to detect potential issues and alerts.

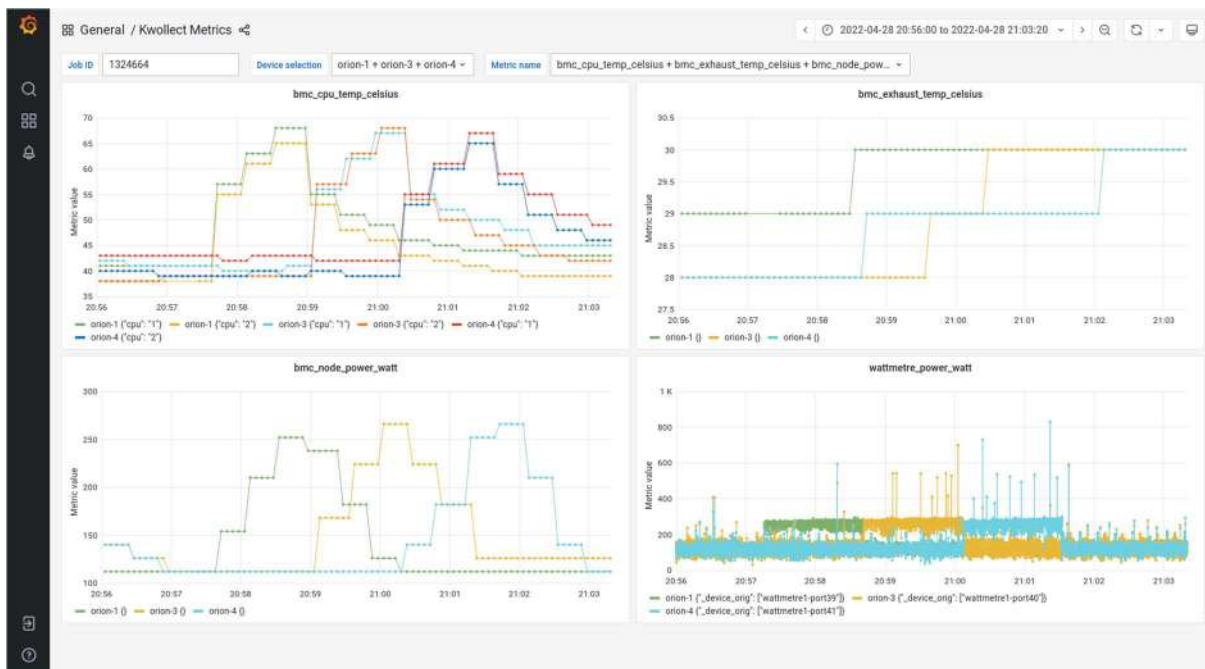


Figure 8. Example of a profiling dashboard provided by Grafana and the Kwollect infrastructure.



The considered capabilities of monitoring concern:

- **Measured metrics:** the considered metrics can cover performance aspects (time, energy consumption, and bandwidth) and infrastructures usage. As SLICES proposes to support with end-to-end experiments, SLICES needs to be able to collect metrics that will be unavoidable and useful for SLICES users.
- **Capturing, collecting, storing & exposing:** Digital equipment (servers, storage and network equipment) include more and more internal sensors, which allow addressing some metrics. Some additional hardware facilities can also be added in order to include external metrics (like energy consumption, or external temperature). The level of granularity of components (from hardware to software) must be clearly defined and expressed. Quality and frequency of data collection must be explored in order to expose relevant information. The complete workflow for capturing, collecting, storing and exposing monitored and profiling data must be specified. This workflow needs to address scalability and robustness issues.
- **Tools and frameworks:** the eco-system of available monitoring and profiling frameworks and tools is rich and diversified. It includes solutions from historical ones like Munin,¹² Ganglia¹³, Collectd¹⁴, etc. to more recent and modern ones like Prometheus¹⁵, Graphit¹⁶, kwolect [KWOLLECT21] as illustrated in Figure 8. Re-using, adapting and designing adapted tools and frameworks for SLICES will need to encompass the heterogeneity of deployed resources. Such solutions need to be studied from various angles such as hardware compatibility, underlying used technologies, intrusiveness, quality of measurement, quality of documentation, their strengths, and their limitations.

3.4.6. Data Management Service

3.4.6.1. Storage Service

In the storage service WG, we have defined different types of storages, to have a clear terminology.

- **Local service specific storage:** e.g., some infrastructures might have a shared nfs storage that is available only on the hardware of that infrastructure. It is not accessible from outside. As it is service specific, we do not think anything needs to be standardized, apart from the fact it should be usable with the standard account.
- **Storage that is accessible easily from everywhere:** this storage can be used from multiple services and sites. It can be used to store data during experiments and in between experiments. The storage needs to support the SLICES accounts. Multiple such storages (e.g., 1 per node) can exist, but should have the same API. Experimenters can then choose for storage nearby their experiment (for performance) or based on

¹² <http://munin-monitoring.org/>

¹³ <https://github.com/ganglia/>

¹⁴ <https://www.collectd.org/>

¹⁵ <https://prometheus.io/>

¹⁶ <https://graphiteapp.org/>



other policies (e.g., some storage/nodes might not allow sensitive/medical data). For pre-operations, we think that storage with S3 API is ideal for this.

- **Versioned storage:** when combining storage of data with metadata about this data, a much richer storage service can be created. See also Data Management service (DMI) below.
- **Artefact storage:** When publishing, it is handy to have an artefact storage for storing long term data with a known URL/DOI¹⁷.
- **Long term storage:** SLICES might also offer a service for long term (>10 years) storage of datasets, etc. with guarantees on availability (e.g., by having multiple copies in multiple nodes/countries).
- **Git server for code:** It might be interesting to have a SLICES git server to store code for long term (e.g., accompanying artefact storage). The advantage is that you see the history of the code (compared to a tar ball as artefact for example).

For operations, SLICES should also look on other storage offerings¹⁸, as those might already offer the right service.

3.4.6.2. Data Management Service

The Data Management Infrastructure (DMI) is designed to support different types of digital objects, such as datasets and machine learning models that are utilized in diverse experimental studies including advanced wireless networking, IoT, artificial intelligence, energy efficiency, or security and privacy. The DMI must support the whole data lifecycle. It should provide interfaces to experiment workflow and staging. The DMI is in development with consideration to specific user needs, which are and not limited to as follows:

1. **Data Storage and Access:** scalable, distributed, accessible and supports wide range of data formats.
2. **Data Integration and Interoperability:** support for integrating various data from different data sources and interoperable with external platforms and systems such as EOSC.
3. **Data Management:** incorporates metadata management with Metadata Registry Service (MRS), data cataloguing to facilitate easy search and retrieval of data and, version control to track changes in the datasets and versions to maintain data integrity and reproducibility.
4. **Data Security and Privacy:** support (trusted) data exchange and transfer protocols that allow policy-based access control to comply with the data protection regulations.
5. **Data Analysis and Processing:** automation of experiment workflows, data pipelines (cleaning, transformation, and analysis) and support for integrating complex ML/AI algorithms.
6. **User and Application access control with identity management:** policies adopted by SLICES community that can be potentially federated with the EOSC Federated AAI.

¹⁷ <https://ckan.iotlab.com/organization/slices-sc>, Last accessed: 2024-06-06.

¹⁸ <https://eudat.eu>, Last accessed: 2024-06-06.

Other specific user needs are identified as quality assurance and regular monitoring of data and DMI as a whole, data backup and recovery protocols, and flexibility to include changes in research and experiment requirements and advancement in existing technologies.

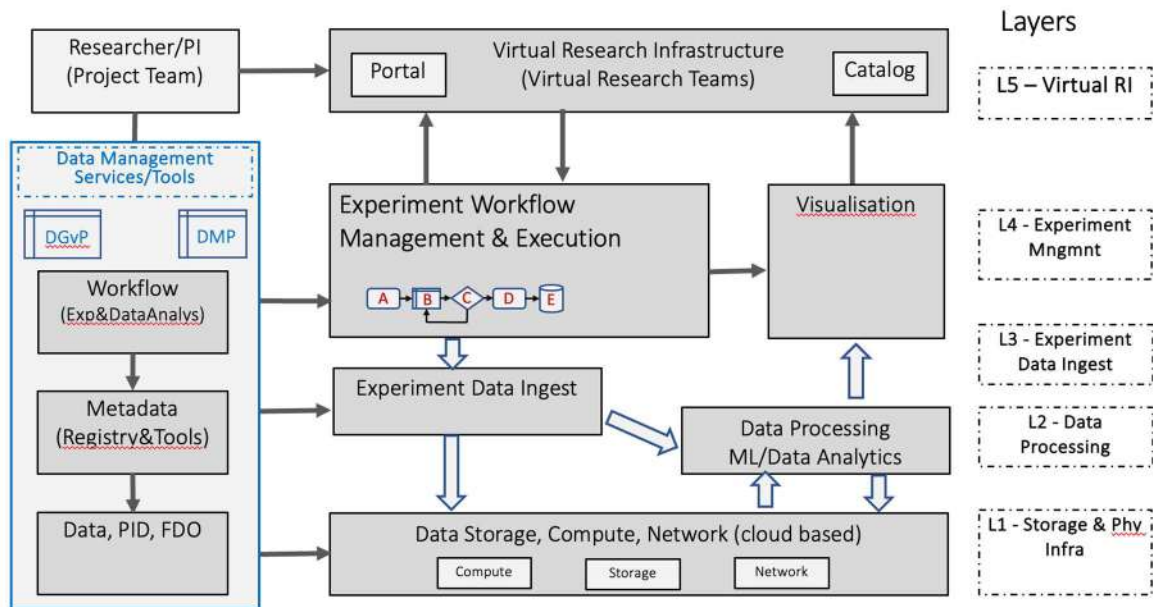


Figure 9. Data Management Infrastructure Architecture.

The DMI architecture, illustrated in Figure 9, consists of five different layers. All layers support the different stages of the data life cycle during an experiment, what includes data storage, data processing, data ingest, data management and data visualization. For the full support of the experiment life cycle, it is essential that every step of the experimental process diligently stores all utilized or generated digital objects and associated metadata to make them available for replication and reproduction. Furthermore, some experiments may require interoperability with external platforms and systems, such as EOSC, e.g., for utilizing external datasets.

To this end, SLICES offers a versatile Experimental Research Data Management Infrastructure that employs end-users with sophisticated data management services and tools to manage and execute complex experiments, harvesting all the digital objects utilized by experiments. Within the DMI, experiments are considered as Digital Research Objects, described with appropriate metadata to be Findable, Accessible, Interoperable and Reusable. These metadata consist of core characteristics, such as a Persistent Unique Identifier (PID), creators, a description, and related or required objects, and experiment-specific metadata, such as the full experiment setup and configuration, experimental workflow, inputs and outputs.

The experiment workflow design is based on the Plain Orchestrating Service (POS), described in Section 3.4.7, developed by Technical University of Munich (TUM) and extended to support the full research data management life cycle, including the stages of:

- 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.

The experiment workflow is supported by a generic information model geared toward reproducibility, with appropriate metadata defined and applied for all experiment components and stage.

SLICES implements a flexible hierarchical metadata model, named SLICES FAIR Digital Object (SFDO), combining domain-agnostic metadata for easy and uniform discovery, and type- and domain-specific information for enhancing machine-actionability, allowing internal and external services to access more complex information about the object and take appropriate actions. SLICES’s hierarchical metadata profile design ensures the full support of FAIR principles at the top level of the hierarchy where carefully selected machine-readable metadata attributes allow for the easy discovery of data by humans and computers. The second and third levels improve machine- and human-understandability and machine-actionability, allowing services to access more information and understand complex and domain-specific metadata structures to take appropriate actions. The model’s architecture consisting of the above three levels is illustrated in Figure 10.

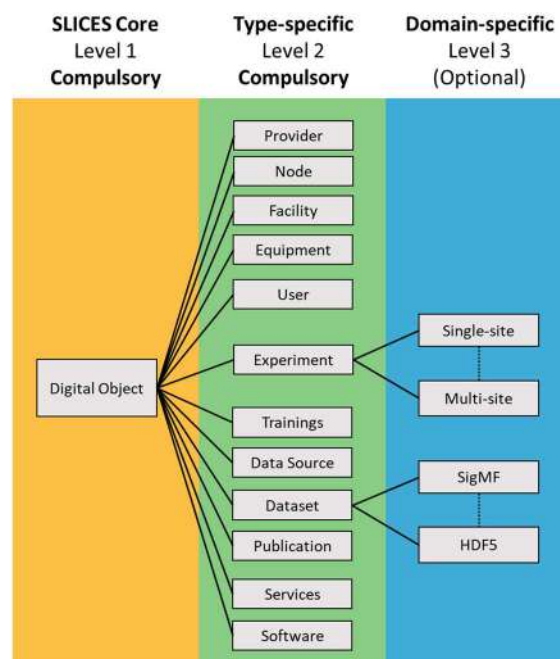


Figure 10. SLICES Hierarchical FAIR Metadata Model (SFDO).



In particular, the first level consists of compulsory domain-agnostic information that can describe any digital object (e.g., data, services, publications, tools), ensuring that it conforms to FAIR principles and beyond. It includes basic information, such as identification, description and its resource type. Management information is also included such as version and metadata profile used. Additionally, SLICES employs a second level of compulsory metadata attributes that are type-specific to enhance machine-actionability for specific commonly used types of digital objects, such as data, services, and software. For example, a dataset may have start and end dates, a facility may have an address. Finally, the third level incorporates optional domain-specific attributes to further enhance interoperability for specific communities. For example, the SigMF standard is designed to record signal (e.g. wireless radio transmissions) data and includes some predefined metadata attributes, such as `core:sample_rate`, `core:datatype` and `core:hw` (hardware that was used for signal recording). This information is exposed as Level 3 SFDO attributes to facilitate enhanced discoverability across different such datasets.

In addition, the project is developing and deploying new tools to further support the researchers, such as the metadata registry service that implements the SLICES metadata model and provides sophisticated metadata management tools to experimenters, such as simple and advanced search capabilities, reporting dashboards, usage and monitoring information, and management tools.

Finally, the project is monitoring closely new developments in the field and participates in multiple support actions with the objective to enhance:

- Interoperability with other models, such as ROcrate and Signposting;
- Interoperability with other platforms, such as the EOSC catalogue;
- Machine actionability and machine actionable data management plans;
- Alignment with the EOSC metadata system and crosswalk registry.

3.4.7. Experiment control and orchestration

Reproducible experiment workflows: Reproducibility is a topic that is often neglected in computer science as it involves additional effort by researchers to test, document, and prepare experimental artefacts to be reproduced by other scientists. However, initiatives, such as the badges for successfully reproduced papers¹⁹, show that this topic is gaining attention in our community. The plain orchestrating service (pos) [CONEXT21] is a framework and methodology for creating and executing reproducible experiments in SLICES. The reproducibility is a core feature of the SLICES/pos experiment workflow. We ensure the creation of reproducible experiments through a well-defined experiment workflow, i.e., *reproducibility-as-a-service* (RaaS). If users adhere to the provided experimental template, the resulting experiments will be inherently reproducible. This way, the effort of creating reproducible experiments is shifted from the experimenter to the testbed framework.

¹⁹ <https://www.acm.org/publications/policies/artifact-review-and-badging-current>, Last accessed: 2024-05-16.



Ongoing developments of the SLICES/pos controller: In the past, the pos workflow was only supported on pos-orchestrated testbeds and limited to our own testbeds at TUM. However, we successfully demonstrated [WONS24] that the workflow can be ported to other testbeds, such as the Chameleon or CloudLab testbeds. To improve the reuse of experimental results, we are currently evaluating different possibilities. A promising standard for packaging, archiving, and exchanging experimental results, we selected the RO-Crate standard²⁰. This is a packaging format that includes metadata in the JSON-LD format to describe various types of experimental results, e.g., data or publications, in an exchangeable format. Our publication [IFIP24] describes how this format can be integrated into the SLICES/pos workflow. We plan to further integrate this standard in future versions of pos. This way experimental results will be automatically packaged in the correct format without additional effort by the experimenter.

Integration of pos in the SLICES pre-operation: We are currently planning, implementing, and deploying the first version of pos for the SLICES pre-operation. The service will be hosted on a server in our testbed. The installed pos framework is currently adapted to cooperate with and integrate other SLICES services described in this deliverable. To demonstrate the integration of these services, we are currently setting up an experiment that demonstrates the collaboration of different services and combines them into a consistent experiment workflow. Two services that we want to highlight are the authentication portal created and hosted by IMEC and the post-5G Blueprint. The post-5G blueprint will include distributed testbed sites such as IMEC in Belgium (for the SLICES user authentication service and the post-5G core), Sophia-node in Sophia Antipolis (for the RAN part of the blueprint), and TUM in Munich (for the deployment of the pos framework). The users will also be able to use the previously described metadata registry service (MRS) to record the metadata of the executed experiments.

3.5. SLICES Blueprint Services

This section presents in more depth the services associated to three initial research domains that have been identified in Section 1.2. These services are expected to rely for their implementation of services provided by Supporting and Basic Services.

3.5.1. Post-5G Blueprint

Data communication is everywhere and cellular networks became the standard. As such, most fields of research rely on cellular communication, either explicitly or implicitly. Moreover, 5G specifications (ETSI TS 123 501) finally made cellular network modular and easily adaptable, broadening the scope of researchers that can actually work on the technology and try new idea. But cellular networks became extremely complex nowadays (and expensive to own and operate) as they have strong bound between hardware, software, cloud, and networking.

The vision of the post-5G blueprint is to allow everyone to take benefit of post-5G in their research, from the “5G agnostic” researchers to the hardcore 6G researchers. To reach this

²⁰ <https://www.researchobject.org/ro-crate/>, Last accessed: 2024-05-16.



objective, the post-5G service offers at the same time a comprehensive heterogeneous post-5G infrastructure distributed over multiple sites and region plus a complete set of software to support the following 5 classes of experiments.

- **Vertical service integration and testing:** this class is proposed to support experimenters that want to benefit from post-5G networks in their research. Fundamentally the research in this class is not related to post-5G but it relies on it. The service is designed such that researchers can run their experiments using all the goodies of a cutting-edge post-5G network without having to setup or know anything about 5G.
- **Software Defined Networking:** this class aims at offering access to the internals of a post-5G network to researchers that work on optimizing the behavior of the network. Most of the time, such research consists in algorithms (e.g., written in xAPPs) that read metrics from the network and automatically tune the network settings to reach some optimization objective. In this class, researchers benefit of real complex networks to try their new findings without the need of deploying or owning them; only what is really needed for them is exposed.
- **Radio/network development:** In this class, researchers want to change the internals of the network by changing its protocols. Researchers can change the implementation of the network and the radio without having to physically own and operate the network.
- **Low-level access to radio resources:** in this class researchers change physical resources to include their own ones. Typically, researchers come with a new physical equipment (e.g., a meta-surface, a THz terminal) and integrate it in the infrastructure to test how it behaves in a true system. This class is specific as it requires (temporary) changes in the physical infrastructure and is thus of exclusive usage by nature.
- **Joint use of post-5G infrastructure and HPC resources:** this last class aims at offering the ability to see post-5G and HPC resources as a continuum. As physical resources and radio in particular are scarce by nature, it makes sense to build digital twins of the infrastructure and link the digital twin with the real infrastructure. Nevertheless, this class is not limited to that. Empowering HPC resources and connect them directly to a telco infrastructure allows to automatically generate code and tests and confront them with real infrastructure in a single unified loop, a bit like a CI/CD with hardware. This is particularly interesting for cybersecurity research with intensive need of compute resources and access to hardware (e.g., for fuzzing).

The post-5G service is designed for experiments, not as a testbed. As such, all actions are defined as experiments that follow the lifecycle presented in the above sections, including transparent collection of data and metadata and their publication within the DMI.

3.5.2. Edge/Cloud Blueprint

In the mid-to-long term, the cloud continuum (CC) service(s) aims to offer practitioners and researchers a playground for assessing algorithmic techniques and mechanisms, involving a continuum of heterogeneous resources, including compute, networking and storage. Rather than a single service, it can be viewed as an entry point where essential building blocks aiding



experimentation into CC scenarios are available and ready to use. Considering the decentralized nature of the edge, and its proximity to the sources/actuators, the following are some representative classes of experiments that are envisaged:

- **(Joint) Resource management:** considering that small-scale compute & storage can be deployed near cellular tower sites, homes, offices, public infrastructure or factory floors etc., there is a potential for microservice-based applications to be delivered in proximity to the user. In this setting, several hierarchical layers of edge nodes, with different capabilities, can be deployed, distributing the resources along to support the execution of applications and their data storage, giving rise to a more fluid edge network model. A seamless integration of all the levels of the infrastructure and novel management approaches that coordinate and orchestrate all its virtualized resources, vertically and horizontally, while ensuring QoS, are paramount. Herein algorithms and mechanisms for resource management across edges and in the edge-cloud, scaling the computation with the data, preserving the applications' Service Level Agreements (SLAs), are one area of investigation.
- **Data fusion/processing:** under this category fall experiments dealing with scalable algorithms for (big) data fusion/processing leveraging a geographically distributed platform for their evaluation & benchmarking.
- **Network resource management:** hardware & software acceleration techniques, in-network processing and time-sensitive flow scheduling algorithms are some relevant areas of investigation that fall under this category.
- **Machine learning:** under this category fall experiments servicing the design of advanced AI solutions blending centralized and decentralized techniques. To design such solutions, advanced services are needed to support efficient transfer, storage and indexing of data, and (potentially) network configuration across a continuum of devices with varying resource constraints where AI modules will be executed both at training and inference phases.

It is clear the services needed to support the experiments vary and cannot be qualified as strictly belonging to the CC service(s) alone. Also, the on-boarding of the experiments and the degrees of freedom experimenters have within a specific testbed is scenario specific. The CC service(s) have the ambition to serve and onboard a wide range of communities conducting research in the IoT/Edge/Cloud domains, while also serving specific needs of verticals in terms of security/privacy, experiment reproducibility and more. To this end, methodologies for infrastructure/experiment provisioning & control and distributed control of resources and data have been identified as paramount. The objective is to adopt & extend current frameworks to easily support service (resource) management over heterogeneous cloud-edge providers, an essential service for the CC blueprint. The point is also to support this at a sufficiently high level of abstraction that is not too hard for experimenters to describe the distributed deployment and its requirements for the controller.

3.5.3. Federated Learning Blueprint

The basic infrastructure service offers bare metal servers (with GPUs) and virtual machines, but most data scientists are not interested in setting up all drivers, CUDA software, etc (and



fighting the version incompatibilities). Hence a higher layer service for machine learning and federated learning (using multiple locations) is needed.

This is exactly what the federated learning service wants to offer: an easy setup for machine learning and more especially federated learning experiments.

We have identified multiple use cases based on interaction with research communities, e.g.:

- Machine learning based intrusion detection for industrial IoT in a multi-tier architecture;
- Scalable fleet monitoring and visualization for Smart Machine Maintenance and industrial IoT applications (e.g. wind turbines);
- Federated learning for cyber security: SOC collaboration for malicious URL detection.

We have started focussing on the first use case (as the other 2 and multiple other use cases are spinoffs of this). The Figure 11 below shows the needed setup for this type of experiment:

- Warehouses with industrial IoT devices where intrusion detection needs to be done;
- Clients which generate load/a DDoS;
- The needed bandwidth and separation for the connection between clients and IoT;
- Multiple or distributed clouds to manage the intrusion detection data;
- An advanced machine learning service (e.g. offering GPUs).

The experimenter wants to vary the processing complexity and locations to see what solution works best.

Figure 11 illustrated the multiple underlying SLICES Services needed to offer such an integrated federated learning service.

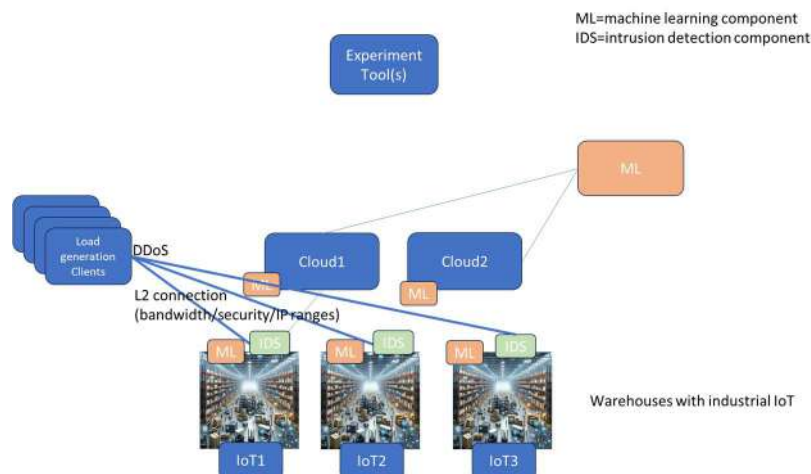


Figure 11. Federated learning service rely on multiple underlying SLICES services.

3.6. Next update cycle



This deliverable presents the view at this stage of SLICES-PP. We will continue to work on analysis new services that SLICES may provide. This section describes a new service that we are just starting to analysis.

3.6.1. The QKD/PQC technology blueprint

There is a great amount of activity related to the irruption of quantum technologies in the IT world, both on the research and operational activities. On the one side, these technologies present a midterm future in which quantum computing will be able to break some of the more widely used public key cryptographic algorithms, but on the other side it brings the possibility to bring information-theoretic security (also called unconditional security) with Quantum Key Distribution (QKD).

Meanwhile as a way to fill the security gap, between the actual situation and the future, Post Quantum Cryptography algorithms (key exchange and signature algorithms) are being developed.

As result, while quantum-based key distribution is steadily being deployed in some scenarios, more or less related to the EuroQCI, national and local QCIs. Access to these resources is not widely available.

This blueprint intends to bridge the divide and offer tools to experiment with Quantum/PostQuantum technologies:

- QKD;
- PQC;
- Quantum Networks;
- Quantum Computing (not directly addressed).

At the technology level, we will address:

- QKD technologies;
- PQC evaluation;
- Data/Quantum coexistence on optical networks.

At the user level:

- Integration of PQC/QKD technologies in applications (including verticals);
- Integration in existing communications services.

An emulated quantum link will be used to support at this time long distance experiments while real equipment will be used where possible.

4. Training activities for users

The main objective related to training activities in SLICES-PP is to organize and deliver training events for SLICES-PP with the goal of promoting SLICES-RI research and fostering the



engagement of external R&D community on experimentation in digital science. The training events target both novice and expert users from multiple disciplines and aim to familiarize them with the tools and infrastructure of SLICES-RI. In detail, the technical objectives set by this task are the following:

- To efficiently identify the training needs and training methodologies that will be followed;
- To develop and provide the respective training material to organize SLICES-RI training events (training sessions, webinars, plugfests, hackathons) as well as shared teaching material used in SLICES-RI for teaching basic skills at the convergence of computing and networking (SLICES Academy);
- To provide guidelines, inter-site collaboration incentives and alignment with national programs pertaining to the teaching of key technical skills in the areas of interest of SLICES-RI (SLICES Academy);
- To facilitate researcher mobility, among the SLICES-RI member institutions and for the research community at large, for the exchange of know-how among the users of the facilities.

4.1. SLICES Academy

SLICES Academy (SA) is the generic name given to the Training activities and their access portal currently put forth by SLICES-RI. SA targets multiple intended audiences, in particular:

- Users from the research community;
- Students (both undergraduate and graduate);
- Technical staff (e.g. system engineers);
- Industry;
- General public at large.

It began as an output of SLICES-SC WP4 which has laid the basic foundation for structuring training and provided the initial implementation of the access portal. It includes preliminary content contributions from SLICES-RI members. SLICES-PP will take over the implementation of the SA upon completion of SLICES-SC in September 2024. We describe here SA's main objectives and defer to the upcoming deliverable D4.221 of SLICES-SC for details regarding the structure of the portal²² and initial online content availability.

4.1.1. Main objectives of the SLICES Academy

SA aims to address the need for the development of a high-level digital education ecosystem at the confluence of advanced networking and computing systems and technologies. EU governments have identified various skills gaps²³ that need to be addressed in order to ensure competitiveness on a global scale. To address this, SA aims to organise new curricula for

²¹ SLICES-SC, Deliverable D4.2 "Final Report on training activities and related material", August 2024. To appear.

²² <https://www.slices-ri.eu/slices-academy/>. Last accessed: 2024-06-06.

²³ [https://data.europa.eu/en/publications/datastories/towards-2023-european-year-skills#:~:text=The%20European%20skills%20gap&text=When%20it%20comes%20to%20digital,digital%20skills%20\(Figure%201\)](https://data.europa.eu/en/publications/datastories/towards-2023-european-year-skills#:~:text=The%20European%20skills%20gap&text=When%20it%20comes%20to%20digital,digital%20skills%20(Figure%201)), Last accessed: 2024-06-06.



enhancing top-level researchers' competences and skills in advanced digital systems and networks, in both engineering schools and research centers as well as with major EU industry via continuing education. Although primarily targeting researchers starting at the PhD level, the effort will also benefit masters-level students. Although governments are expressing high concern for closing gaps in basic digital skills, the same can be said for the advanced skills needed by EU operators, equipment vendors and cloud infrastructure providers who have difficulty finding qualified staff for next-generation product development. The skills addressed by SA target the development of modern hardware and software as well as innovative and novel ideas allowing maintaining a significant share of standards-essential intellectual property (patents). Remaining dominant on a global scale requires rejuvenation of the EU workforce in all areas and most importantly increasing the dwindling numbers of students drawn to high-level work in modern networking technologies.

The current skill-set identified by SA are summarized as:

- **Fixed Networking:** software-defined switching, time-sensitive networking, Data-plane acceleration, Cloud-native Networking, Telco security, Quantum networks.
- **Edge/Cloud:** virtualization, containerization, devops, CI/CD, orchestration, far-edge resource orchestration, Cloud Security.
- **Wireless Networking:** OpenRAN, SDR, fronthaul, near RT-RIC/xAPPs, NTN, B5G/6G RAN, cloud-native RAN/Core, wireless networking drivers, IoT devices and energy-efficient protocols, Applications and Services, Optical Wireless, mmWave/THz, Smart Antenna Tech, Positioning, URLLC, mobility, Wireless network security.
- **Compute:** HPC, Quantum, Accelerators (GPU/FPGA), new programming languages
- **AI/ML:** Data analysis, Federated learning for networking, AI/ML for radio-access and core network control.

More refinement based on the needs of the SLICES-RI ecosystem will occur continuously. SA aims to provide an open and mutualized curricula across SLICES-RI institutions. It is the hope that through mutualization of know-how and tools via SA, SLICES-RI can provide concrete impact in closing skills gaps in digital systems in the EU and at the same time modernize the teaching of modern computing for networks in our schools and research centers.

4.1.2. Licensing of Pedagogical Content

An important aspect of SA is the licensing of content in an open-curricula system. Several options exist and a common framework for licensing SLICES-RI pedagogical content should be agreed upon by the SLICES-RI members. This will allow for collaborative development of content across institutions and through mutualization will facilitate use of up-to-date information. The de facto licensing strategy is the Creative Commons²⁴ family (CC) of licenses which are summarized here. These are used by the multimedia repository of Wikipedia and many other content systems. CC comprises four types of licenses:

²⁴ <https://creativecommons.org/licenses/>, Last accessed: 2024-06-06.



- **Attribution (CC BY):** This license lets others distribute, remix, adapt, and build upon your work, even commercially, as long as they credit you for the original creation. This is the most accommodating of licenses offered. It is recommended for maximum dissemination and use of licensed materials.
- **Attribution-ShareAlike (CC BY-SA):** This license lets others remix, adapt, and build upon your work even for commercial purposes, as long as they credit you and license their new creations under the identical terms. This license is often compared to “copyleft” free and open-source software licenses. All new works based on yours will carry the same license, so any derivatives will also allow commercial use. This is the license used by Wikipedia, and is recommended for materials that would benefit from incorporating content from Wikipedia and similarly licensed projects.
- **Attribution-NoDerivs (CC BY-ND):** This license lets others reuse the work for any purpose, including commercially; however, it cannot be shared with others in adapted form, and credit must be provided to you.
- **Attribution-NonCommercial (CC BY-NC):** This license lets others remix, adapt, and build upon your work non-commercially, and although their new works must also acknowledge you and be non-commercial, they don’t have to license their derivative works on the same terms.

Content distributed as part of software (e.g. markdown or .md files) can also be licensed along with the corresponding software license such as GPLv3, Apache, BSD-3, MIT, etc.

SLICES-RI needs to define a strategy for content distribution in parallel to the establishment of the central office. One of the key items to decide among the partners is the copyright of the content in the context of SLICES-RI. The simplest possible strategy is to request SLICES-RI partners to donate content to the SLICES-RI ERIC while retaining copyright. This is a similar mechanism to licensing software with foundations (e.g. Linux Foundation). Contributors to content on SLICES-RI repository could also potentially agree to a “contributors license agreement” (CLA) put forth by the SLICES-RI ERIC which stipulates the use of the content prior to actually contributing to a SLICES-RI content repository (e.g. GitLab or SA). As described on Civic Commons²⁵, “The contributor retains formal copyright ownership of the contribution, but gives the project (SLICES-RI) a (deep breath) non-exclusive, perpetual, world-wide, irrevocable, no-charge, royalty-free copyright license to use, reproduce, prepare derivative works of, publicly display, publicly perform, sublicense, and distribute (etc.) the contribution. In other words, the contributor retains copyright, but promises not to exercise most of the powers that copyright ownership would ordinarily imply, thus enabling the project (SLICES-RI) to feel safe in accepting the contribution”. SLICES-RI can keep track of contributions that it distributes and thus its use by members for pedagogical purposes. It can also be forked and distributed elsewhere. Content managed in this way is collaborative by nature, so that it can be built by the community. Moreover, contributions from the community-at-large (i.e. not necessarily SLICES-RI member institutions) can be accepted if the CLA is signed by contributors. This will additionally help put make the SLICES-RI pedagogical offering available internationally.

²⁵ https://wiki.civiccommons.org/Contributor_Agreements/, Last accessed: 2024-06-06.



At the stage of this deliverable, we recommend:

- CC BY or CC BY-SA should be used for SLICES-RI content. The primary advantage of CC BY-SA is incorporation of Wikipedia content in SLICES-RI content with the hope that the opposite can also happen in time. CC BY is enough to publish on YouTube, Internet Archive, Vimeo, Wikimedia Commons.
- For SLICES-RI content, prior to an agreed policy upon creation of the SLICES-RI ERIC, institutions should use CC BY or CC BY-SA for the time being.
- Discussion on the use of a CLA by the SLICES-RI ERIC should take place now and SLICES-RI member institutions should be made aware of this possibility to avoid difficulties along the way that should not hinder scientists' contributions to SLICES Academy or other SLICES-RI content.

4.1.3. Alignment with national initiatives

National governments are also mobilizing efforts to rejuvenate pedagogical offerings in digital systems and networks at all levels of study. The underlying programs will provide significant funding for achieving the same goals as SA and can be used to mutualize efforts with SA for partners willing to do so. For instance in France the programme entitled "Compétences et Métiers d'Avenir" which is part of France 2030 national funding has granted several multi-year projects comprising public and private institutions²⁶ for innovating teaching methods in the areas of interest of SLICES-RI (5G, Cloud, AI, Quantum). Several of these projects may have an overlap with SLICES-RI participants who should strive to mutualize efforts as much as possible with SA. One identified project is IMTFor5G+ which includes schools from IMT and EURECOM that are involved in SLICES-RI. Explicit links to the French node (SLICES-FR) have been made and EURECOM is committed to mutualizing pedagogical material and platforms for SA content from its SLICES-RI site in the context of IMTFor5G+. EURECOM is working towards this harmonization goal with other IMT partners active in SLICES-RI (e.g. Telecom Paris).

At this stage, SLICES-RI is in the process of identifying similar funding instruments in other national nodes that can be exploited for developing SLICES Academy content.

4.2. Organization of Events

The primary events organized by SLICES-RI are:

- **Summer schools**, which act as multi-disciplinary opportunities to train new users of the RI. The summer schools will continue to help interested students and academics as well as individuals working in academia, research or the private sector, to better understand the specific opportunities of the RI within the school's selected thematic priorities and get comprehensive knowledge on the various aspects of the selected theme. So far, two summer schools have been organized by SLICES-RI (Volos Greece

²⁶ <https://www.info.gouv.fr/actualite/france-2030-annonce-de-9-laureats-competences-et-metiers-d-avenir-en-faveur-de-la-souverainete>, Last accessed: 2024-06-06.



2022, Oulu Finland 2023) and a third will occur in 2024 in Lipari Italy. SLICES-PP will continue to propose summer schools starting in 2025. The number of schools (winter/spring) should also reflect the increase in subject areas as SLICES-RI grows.

- **Training at major international conference venues.** In general, these training sessions expose technologies put forth by SLICES-RI to an international audience and allow for exchange with similar experimental facility initiatives internationally. Multiple SLICES-RI related training sessions should occur each year at major IEEE or ACM conferences or workshops. Given the increased scope of SLICES-RI as of 2025, the RI should aim for one such event in key areas (next generation networking, Cloud/Edge Continuum, Open Data Management) at major international conference or workshop venues. For example, in 2023 a tutorial was organized at the IEEE Conference on High-Speed Switching and Routing and in June 2023 on a V1 of the SLICES post-5G blueprint, a similar tutorial/training was organized at IEEE IFIP 2024.
- **Hack-a-thons** for engaging community-development of basic hardware and software tools used by SLICES-RI and related international experimental infrastructure initiatives. SLICES-RI has co-sponsored such events related to the post-5G blueprint and extensions to OpenAirInterface to support software-defined controllers (O-RAN xApps) but also refining the post-5G blueprint reference implementation to support authentication via the SLICES portal. These events allow researchers to collaborate on extending existing tools and to transfer knowledge among tool developers. It is a very explicit way of developing a broad community for mutualizing tools.
- Interconnecting SLICES-RI sites with global open-networking initiatives through **plugfests**. Many communities related to the Linux Networking Foundation organize plugfests to integrate technologies and software tools for experimentation (e.g., OPNFV, Open Compute Project, Nephio, O-RAN). SLICES-RI will periodically offer its sites for such events along with specific training on their use and interconnection capabilities. These are opportunities to interact with industry-focused groups and allow for transfer to and from SLICES-RI tools and technologies to impact the area at large. SLICES-RI had interactions in this direction with O-RAN plugfests using the SophiaNode site in Sophia Antipolis.
- SLICES-SC has continued the organization of “The Networking Channel” series of **Webinars**. These regular thematic events help disseminate SLICES-RI research results and bring international experts in the field alongside SLICES-RI. They are very valuable for EU researchers, in particular PhD students, and have enjoyed a very high attendance internationally. Such methods should be encouraged in other areas of SLICES-RI.

In addition to these existing efforts, to cover the span of subjects’ areas put forth by SLICES-RI, additional events will need to be organized at all levels, national, EU and international. It is envisaged that the number of thematic schools associated with SLICES-RI will increase. The current offering was limited primarily to the partners in SLICES-SC and DS and only covered a subset of the areas of interest. Given the expected duration of SLICES-RI and its number of prestigious academic research and teaching organizations, it would also be interesting to consider founding and hosting a recurring annual international workshop (IEEE, ACM) on the subject of experimental computing and communication technologies.



4.3. Mobility of Researchers

Researcher mobility is an essential component of sustaining a multi-site/national experimental facility. The goal of this activity is to increase links for enabling transnational research in digital sciences and to promote the knowledge flows and collaboration between institutions promoting new interdisciplinary academic-industrial networks. This activity is targeted both to members of the consortium, as well as researchers not affiliated with any of the consortium partners. This mobility initiative aims to transfer skills and expertise between academia and industry in order to contribute to the development and the improvement of a common understanding of research infrastructures for Digital Sciences (and of SLICES-RI in particular) and to foster cooperation between stakeholders.



Bibliography

[CONEXT21] Sebastian Gallenmüller, Dominik Scholz, Henning Stubbe, Georg Carle, “The pos Framework: A Methodology and Toolchain for Reproducible Network Experiments”. In Proceedings of the 17th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '21). Association for Computing Machinery, New York, NY, USA, 259–266. DOI: 10.1145/3485983.3494841

[DS-D1.1] SLICES-DS, Deliverable D1.1 “Technological status and capabilities of existing ICT Research Infrastructures”, March 2021. Available at URL http://www.slices-ds.eu/wp-content/uploads/2021/12/SLICES-DS_D1.1.pdf.

[DS-D2.2] SLICES-DS, Deliverable D2.2 “SLICES as a Service, baseline”, September 2021. Available at URL http://www.slices-ds.eu/wp-content/uploads/2021/12/SLICES-DS_D2.2.pdf.

[EC16] European Commission, “European Charter for Access to Research Infrastructures, Principles and Guidelines for Access and Related Services”, 2016, DOI: 10.2777/524573.

[ESFRI20] ESFRI Whitepaper 2020, “Equipping Europe with infrastructures for groundbreaking research ESFRI's response to the new challenges lying ahead for Europe.” Available at URL https://www.esfri.eu/sites/default/files/White_paper_ESFRI-final.pdf.

[IFIP24] Eric Hauser, Sebastian Gallenmüller, Georg Carle, “RO-Crate for Testbeds : Automated Packaging of Experimental Results”, SLICES Workshop @ IFIP Networking, June 2024.

[KWOLLECT21] Simon Delamare, Lucas Nussbaum, “Kwollect: Metrics Collection for Experiments at Scale. CNERT 2021 - Workshop on Computer and Networking Experimental Research using Testbeds”, Virtual Event, May 2021.

[WONS24] Henning Stubbe, Sebastian Gallenmüller, Georg Carle, “The pos Experiment Controller: Reproducible & Portable Network Experiments”, 19th Wireless On-Demand Network Systems and Services Conference (WONS), Chamonix, France, January 2024, pp. 1–8.

