Ref. Ares(2021)5478711 - 06/09/2021

slices DS

HORIZON 2020
H2020 - INFRADEV-2019-3

D2.2 SLICES as a Service, baseline

Acronym	SLICES-DS
Project Title	Scientific Large-scale Infrastructure for Computing/Communication Experimental Studies – Design Study
Grand Agreement	951850
Project Duration	24 Months (01/09/2020 – 31/08/2022)
Due Date	31 August 2021 (M12)
Submission Date	6 September 2021 (M13)
Authors	Christian Perez (Inria), Luke Bertot (Inria), Walid Dabbous (Inria), Laurent Lefevre (Inria), Nathalie Mitton (Inria), Lucas Nussbaum (Inria), Adrien Lebre (IMT), Nikos Makris (UTH), Cedric Crettaz (MI), Bartosz Belter (PSNC), Carmen Guerrero (UC3M), Panayiotis Andreou (UCLAN), Nearchos Paspallis (UCLAN), Andreas Pamboris (UCLAN), Josephine Antoniou (UCLAN), Irene Polycarpou (UCLAN), Yuri Demchenko (UVA), Kishor Joshi (UVA).
Reviewers	Serge Fdida (SU)

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951850. The information, documentation and figures available in this deliverable, is written by the SLICES-DS 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.

Scientific Large-scale Infrastructure for Computing Communication Experimental Studies **Design Study**

www.slices-ds.eu

Executive Summary

This document provides a description of the baseline services provided by SLICES to experimenters based on their various demands and needs as well on access modes and types. It starts with an overview of Slices mainly based on Deliverable D2.1. Then, it analyzes user requirements for SLICES services based on Deliverable D1.1 results. It also introduces three examples of usages of the platform in the field of cloud/edge, IoT, and wireless/wired networking. Then, the document presents the access modes and types of SLICES. The three types of interfaces offered by SLICES are then described (APIs, Web/Notebook, domain specific interfaces) before presenting and detailing the initial set of twelve services:

1. [USERS_MGT] User and group management
2. [EXP_MGT] Experiment management
3. [DISCOVERY] Resource discovery and description
4. [RESERVATION] Resource reservation
5. [CONFIGURATON] Resource configuration
6. [MONITORING] Resource monitoring and profiling
7. [ORCHESTRATION] Experiment control and orchestration
8. [DATA] Data Management Service
9. [ANALYSIS] Experiment data validation and correlation with other experiments
10. [DASHBOAD] Dashboard
11. [DOCUMENTATION] Documentation and Online Experiment Helpdesk
12. [ACCOUNT] Accountability & billing

The usage of some of these services is illustrated in four basic and typical examples: a basic experiment, an orchestrated experiment, a basic data-intensive experiment, and a basic course.

Table of contents

1 SLICES Overview

SLICES aims at providing high quality experimentation services with emerging technologies around the area of digital sciences. The goal is to build a large-scale infrastructure for experimental research in computer science, and more precisely in networking and distributed systems, targeting scientific challenges in the fields including wireless networking, IoT, edge/fog/cloud computing and distributed systems.

SLICES will be a highly distributed infrastructure, to reflect the fact that the environments we aim to study are themselves distributed (e.g., Fog/Edge computing), towards supporting a large variety of viable topologies in distributed computing systems. SLICES shall not be a light federation of independent sites, but rather a coherent environment to perform large scale distributed experiments. The knowledge and experience gathered from previous efforts and initiatives have resulted in the design of several tools and platforms that shall manage this infrastructure in a coordinated way, providing users with a consistent environment that shall overcome the technical challenges of multi-sites experiments.

As described in Deliverable D2.1 "Initial description of the SLICES architecture" [D2.1], SLICES will provide a fully programmable remotely accessible infrastructure to the Digital Infrastructure scientific community. The respective frameworks are designed and they will be developed for ensuring seamless and easy access to the experimental resources. The different site facilities will form an integrated single pan-European facility, which experimenters will access seamlessly, without noticing (if not needed by specific experimental requirements) the fact that resources might come from different facilities. The integrated facility will adopt common tools for managing and orchestrating experiments over the infrastructure, as well as provide a single access and credentials to users. An initial version of its architecture used for the management is described in Figure 1.

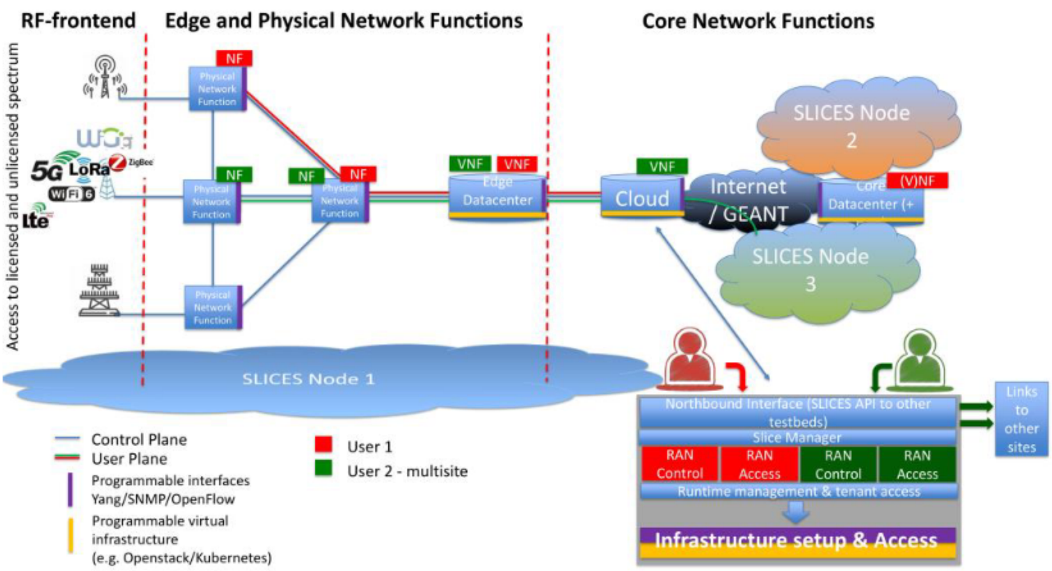

Figure 1: SLICES Architecture Layout Overview

This document deals with the computer-based services envisioned to access and make use of SLICES from a user point of view. It does not cover human-based types of services such as user training sessions or specific support to design or optimize experiments. Of course, training will be a very important element with online documentation and the possibility to set up online tutorials. In addition to general user support, our experience in operating research infrastructures has also shown that some experiments need a specific support to help users in handling their complexity. While SLICES will provide such human-based services on top on the research infrastructure, this document only focuses on computer-based services.

Given the diversity of usage (wireless/wired networking, IoT, Cloud, Distributed systems) and the profile of users (academic/industry), very distinct users with various experiment requirements are expected. Moreover, the duration of an experiment could range from short duration (in the order of minutes) to longer duration (hours) or even to very long duration (a few days).

The remaining of this document is organized as follows. Section 2 presents a need and requirement analysis as well as some examples of usage. Section 3 deals with the access modes and types under which the platform will be provisioned while Section 4 describes the user interfaces of the platform. Section 5 provides an initial list of services facilitating access to the experimental equipment while section 6 illustrates the usage of some of these services with four examples. Section 8 provides a general conclusion and references are given in Section 9.

2 User Requirements: Analysis and Examples

This section first sums up the analysis of user requirements with respect to the services that SLICES has to provide. Then, it presents three particular types of experiments as extracted from user needs and requirements.

2.1 User Requirement Analysis

Towards gathering the needs from all the relevant scientific communities that will benefit from the operation of SLICES, two tools were used to collect user feedback on their needs and requirements. These tools were a user survey, and a SLICES workshop, advertised and open to relevant ICT communities. The tools, methodologies, and extensive results are discussed in detail in Deliverable D1.2 “Requirements and needs of scientific communities from ICT-based Research Infrastructures” [D1.2].

The requirement analysis showed that the future ICT research infrastructure should be able to handle at the same time the distribution of computing (cloud, fog, edge) and the processes running on top of computing, which should be optimized thanks to Artificial Intelligence (AI) and Machine Learning (ML). AI and ML will also play an important role to manage the underlying layers, in particular to improve the performance and the energy efficiency of the wireless and network technologies (SDN/NFV) in the context of 5G networks and beyond. The security, privacy and scalability are very important features to be implemented by the research infrastructure with the possibility to assess them directly in the RI. Another important challenge to be addressed by the future RI is the possibility given to the researchers to configure for themselves the hardware used in the different kinds of experiments. The geographic distribution of RIs is also a goal for the research community as new use cases could be created in a highly distributed environment.

The requirement analysis indicated that the research infrastructure should use the latest communication technologies such as 5G and beyond 5G, in a real environment including components deployed at the cloud, at the fog and at the edge. More realism in the RI is also required by the researchers to ensure experiments closer to the industry; this means that the scalability should be improved compared to the current testbeds. The open access and the configuration of the research infrastructure are very important for the research community. New tools should be incorporated in the new RI like traffic/usage generators and real data sets. These tools should at the end help the development of AI and ML in experiments. SLICES RI will be designed in order to address support the design assumptions of the emerging 6G.

The use cases that emerged from the analysis cover a wide range of applications, but can be separated in two categories: core network and verticals. The core network encompasses the communications technologies and how the interactions between the components interconnected through these technologies can be realized in an efficient way. Use cases associated to verticals (connected vehicles, smart grid, etc.) should be also supported by the research infrastructure. It means that the new SLICES RI should be able to handle at the same time the generic technologies and their applications in the different ICT domains, taking into account the specific requirements and needs of each use case. The most relevant use cases can be combined. For example, AI and ML are applied in the context of IoT and 5G to optimize the collection and the analysis of data. The network layer can also be better managed using AI and ML in the SDN/NFV deployments. Some other proposed use cases are horizontal and concern the energy efficiency, the scalability, the cybersecurity, and the privacy. An indicative list of the main use cases that need to be supported by the SLICES RI are provided below:

- Optimization done by AI and ML: The dedicated experiments should demonstrate how the AI and ML can optimize the decision making.
- 5G and 6G experiments.
- IoT experiments.
- Cloud, fog, edge computing experiments.
- CO2 reduction, energy efficiency, smart grid.
- SDN/NFV experiments.
- C-V2X and autonomous cars.
- Augmented and Virtual Reality (AR/VR).
- Smart cities.
- Cybersecurity and privacy.
- Distributed infrastructure: How to handle a distributed infrastructure? Real experiments are to be set to have a clear vision for the future distributed infrastructure management.
- Large scale network testing and validation.

2.2 Examples of Experiment Needs

Based on the previous analysis, this section highlights three particular types of experiments extracted from user needs and requirements. Of course, many other types of experiments are envisioned combining for example heterogeneous sets of sensors, networks, and datacenter nodes. However, the examples cover the main areas around which SLICES infrastructure investments are planned, and therefore they are overall well representative of the expected use of the RI. Also, note that different, more fundamental and longer-term experiments will be possible via combining the mentioned services. For example, experiments on novel future Internet paradigms can be run by combining in-network processing components (Section 2.2.1) with small scale, resource constrained IoT devices (Section 2.2.2) with latest high-speed wireless access networks (Section 2.2.3). The possibility to experiment also in SLICES, long-term research ideas is particularly important, as this is almost impossible with today's testbeds, at least at the scale and realism that SLICES will provide.

2.2.1 Services for Cloud/Edge and distributed systems Experiments

SLICES aims at providing an experimental environment that is as controllable, configurable, and observable as possible, building on services that were already designed such as for example for the FIT, Grid'5000 and Virtual Wall testbeds: i) bare-metal reconfiguration of hardware to allow experimenters to use their own software stacks (up to the deployment of complex cloud/edge stacks); ii) fine-grained monitoring of performances, energy consumption, network traffic, and environmental conditions (temperature, humidity); iii) network emulation to allow experimenters to instantiate arbitrary network topologies.

The services related to cloud and distributed systems will be made available on several geographical locations to allow for experiments that target distributed environments (multi-clouds, Edge computing) using the real physical networking infrastructure, in addition to network emulation.

SLICES offering will be mainly based on off-the-shelf hardware (for the time being, x86 and ARM servers), providing access to various kinds of technological tradeoffs in terms of CPU, network (both Ethernet SR-IOV and HPC networks such as Infiniband and Omni-Path), storage performance, and GPUs. SLICES will also include more experimental (as of now) hardware, such as the emerging architectures for Cloud/HPC environments (POWER, RISC-V, etc.), emerging GPU vendors, AI specific servers, as well as integrated Edge computing devices. The RI infrastructure will be upgraded over time following the needs of the SLICES scientific community.

2.2.2 Services for IoT Experiments

SLICES will provide several services for the experiments on IoT devices, which will be available from the Internet. The environment dedicated for IoT experiments will offer the needed services to set up experiments related to the various IoT domains and verticals. This includes the full configuration of the IoT nodes, the installation and the deployment of the required software, the monitoring and control of the IoT devices but also of the IoT Gateways such as LoRa gateways. Of course, each experiment will have its own dedicated space on the IoT testbeds with the configuration specified by the researcher owning the experiment. The idea is to offer full programmable devices at each level.

The environment of the IoT experiments will allow the researchers to monitor the network traffic, the energy consumption of IoT nodes and the behavior of the IoT devices. In the same manner, the topology of the IoT networks will be configurable by the researcher permitting a large panel of experiment possibilities in real conditions.

The IoT part of SLICES will be composed by servers offering virtual environments configurable by the researchers with the cutting-edge virtualization technologies. The IoT devices will come from the market and from the research and development centers and will be programmable and configurable by the experimenters. Other components like IoT gateways or edge computing hardware will be available for the IoT experiments. All these elements allow the building of IoT experiments in various real contexts, as demonstrated by the relevant existing facilities on which SLICES builds. The platform will provide a large scope of heterogeneous devices in terms of computing and memory capacities and in terms of radio technologies. Indeed, already available hardware includes light microcontrollers like TI 430, ARM M3 to much more powerful nodes featuring an A8 microcontroller able to run a full Linux OS going through M4, Arduino. As of now, the deployed hardware will already allow the use of the following technologies WiFi, BlueTooth, LoRa, 802.15.4, NBIoT, LTE-M and beyond thanks to the cognitive radio modules. SLICES will provide the ability to the user to design and use their own hardware, being thus modular, flexible and open to further technologies that can pop up by the end of the project. Different and flexible physical network topologies will also be provided with tools able to perform dynamic topology control.

This shall be enabled by offering the ability to deploy on-demand mobile resources, either as IoT devices or as edge-enabled resources. This will be enabled through the use of fully remotely programmable mobile robots. IoT-oriented edge resources of different kinds will also be deployed, being able to run AI at each level. Different parts of the platforms are deployed in diverse environments, completely open (outdoor) to light controlled (indoor in different life environments, public, offices, etc.) to fully controlled (in anechoic rooms).

2.2.3 Services for Wireless/Wired Networking Experiments

SLICES will allow the prototyping and designing of future wireless technology, from enhanced 5G and further. The proposed services will be modular to anticipate the arrival of potential new technologies. SLICES will deploy a set of highly programmable off-the-shelf cells of different scale (macro-, meso- and small-scale) programmable through dedicated APIs (e.g., O-RAN, TR-96). The programmability of the cells will be exposed to experimenters through SLICES services in order to allow experimenters to interact with the equipment in an organized and unified manner. Apart from the off-the-shelf devices, SLICES will deploy fully programmable hardware radios through Software Defined Radios (SDRs) e.g., the Universal Software Radio Peripheral (USRP) devices that allow reliable flexible wireless propagation. By providing highly programmable wireless technologies (e.g., 5G NR and prototype transceivers for new spectrum, such as millimeter-Wave, Terahertz and visible light communication bands), SLICES will allow researchers to test lower physical layers and experimentally evaluate new protocols that will push forward the research in the domain. Programmability for the wireless networks will extend to the entire stack, by providing open-source solutions for the realization of base stations (e.g., the OpenAirInterface platform providing a base station implementation for 5G entirely in software or Open-Source drivers for WiFi experimentation). As such, networks are highly softwarized, experiments regarding network management components such as service placement and slice orchestration in a fully programmable Radio Access Network will be enabled. Such functionality will allow users from different disciplines to deploy, validate and trial their scenarios on top of a real beyond-5G platform that would otherwise be very hard to find. Experiments around this ecosystem that shall be enabled are the following: experiments requiring high CPU consumption, high data rate, mobility support and low latency such as Multiple access Edge Computing (MEC) functions in the context of disaggregated CU-DU base station, AR/VR applications, handover for services using UAV drones and autonomous vehicles. SLICES will also deploy tools, hardware and software able to perform dynamic massive MIMO in different environments, both indoor and outdoor and programmable at all levels of the network but also of the communication stack

Different parts of the platforms will be deployed in diverse environments, completely opened (outdoor) to light controlled (indoor in different life environments, public, offices, etc.) up to fully controlled (in anechoic rooms). SLICES will offer tools to set up, run, and monitor wireless experiments in different environments though unified tools and application interfaces.

For the design of the transport network between sites, optical fiber links from NRENs and GÉANT will be deployed where this will be available (note that SLICES is already supported by GÉANT and different NRENs), being able to interconnect a subset of the base stations installed in SLICES. Through the integration of Software Defined Networking (SDN) equipment (e.g., P4 programmable switches), experiments regarding high-speed switching networks will be made feasible.

3 SLICES Access Types and Modes

3.3 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:

- Trans-national 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;
- Trans-national 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.

3.4 Access Modes

SLICES will be mostly accessed via the Internet in a remote fashion (remote and virtual access modes) by users. 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"
- "Market-driven";
- "Wide".

Following the analysis of Deliverable D1.1 [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 "marked-driven" 20%. These numbers will be further refined and managed by the CMO based on the recommendations of 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.

From a user point of view, accessing SLICES will share some similarities to accessing existing digital infrastructure testbeds though SLICES will offer a larger scale and more diverse infrastructure as well as an advanced monitored and user-controllable infrastructure for researchers in wireless/wired networking, IoT, Cloud, and distributed systems communities.

Finally, SLICES will be registered into the EOSC catalog. The data produced and offered through SLICES will be annotated appropriately, while SLICES will provide EOSC compatible mechanisms to easily access the platform (authentication and authorization) and to manipulate data (EOSC FAIR mechanisms). This insertion into EOSC will be facilitated as several partners of SLICES are currently involved in EOSC, more are expected to become members of the future EOSC association and thus to participate to the EOSC Strategic Research and Innovation Agenda. Deliverable D4.2 "SLICES infrastructure and services integration with EOSC and Open Science (initial proposal)" [D4.2] provides an initial study of these aspects.

4 Interfaces to the SLICES RI

Several methods to access the platform will be provided to the users, based on the level of control that they need to have over the infrastructure components. For the most experienced users, low level APIs will be provided that will allow fine tuning of complex experiments with a low-level access to the different hardware. Such features will allow the users to have full control over the experiment (e.g., setting the modulation of the transmitted wireless signal during the experiment). The APIs will be further used to develop higher-level tools to automatize or simplify experiments for some specific communities.

A second solution will be to provide access to the SLICES infrastructure through solutions that are widely used, like for example notebooks such as Jupyter [JUP]. Such solutions hide some of the low-level details and offer a single interface for defining the environment, reserving the resources across different sites, running the experiment, and finally analyzing the results and managing the output data.

For less experienced users, simple access through a portal web page will be provided. Such access will allow less control of the experiment during the experiment execution, but it will allow fast experiment bootstrapping for novice users. Providing access through such interfaces will enable the quick replication of the experiment across different sites and will foster the capabilities of the infrastructure on reproducing experimental results.

It is worth noticing that the different levels of access will also correspond to finer or coarser grain control over the deployment of experimental resources. More experienced users would be willing to control exactly which experimental resource to use from which facility at which site, while less-experienced users might not even know that their experiment is actually using heterogeneous resources composed out of different facilities spread across Europe. In the latter case, SLICES, through dedicated management components, will automatically assign resources to the users' experiments, in order to optimize the overall utilization of the RI's resources.

5 Description of Functional User Oriented SLICES Services

This section lists an initial set of services for SLICES users, i.e., experimenters, based on the long experience of SLICES members in running smaller and less diverse testbeds. It does not cover internal services needed to operate the platforms though some services can be of course used for it. The initial list identifies 12 types of services that can be listed by following the life cycle view of the usage of the platform:

1. [USERS_MGT] User and group management
2. [EXP_MGT] Experiment management
3. [DISCOVERY] Resource discovery and description
4. [RESERVATION] Resource reservation
5. [CONFIGURATON] Resource configuration
6. [MONITORING] Resource monitoring and profiling
7. [ORCHESTRATION] Experiment control and orchestration
8. [DATA] Data Management Service
    a. Data storage
    b. Data transfer
    c. Data persistent Id
    d. MetaData generator
9. [ANALYSIS] Experiment data validation and correlation with other experiments
10. [DASHBOAD] Dashboard
11. [DOCUMENTATION] Documentation and Online Experiment Helpdesk
12. [ACCOUNT] Accountability & billing

These SLICES services are structured in four main categories:

1. **User and platform management services**

    1.1: [USERS_MGT] User and group management

    1.2: [DOCUMENTATION] Documentation and Online Experiment Helpdesk

    1.3: [ACCOUNT] Accountability & billing

2. **Resource management services**

    2.1: [DISCOVERY] Resource discovery and description

    2.2: [RESERVATION] Resource reservation

    2.3: [CONFIGURATON] Resource configuration

    2.4: [MONITORING] Resource monitoring and profiling

3. **Data oriented services**

    3.1: [DATA] Data Management Service

    3.2: [ANALYSIS] Experiment data validation and correlation with other experiments

4. **Experiment management services**

    4.1: [EXP_MGT] Experiment management

    4.2: [ORCHESTRATION] Experiment control and orchestration

    4.3: [DASHBOAD] Dashboard

The following sections describes the services grouped by category.

5.5 User and platform management services

5.5.1 User and Group Management

A first service is the access to the platform. Several types of accesses are envisioned, such as a web portal and an API as described in Section 4. The web portal provides an interactive and intuitive access to the platform, in particular for newcomers. On contrary, the API, typically a RESTful API, is mainly dedicated for experimentation automatization, and hence it targets tool developers and advanced users.

A user generally is a member of one or several logical groups such as institutions, company, scientific communities, consortia, etc. Hence, this service shall be able to provide an interface to let a user select her/his profile as well as provide for authorized users to create new groups and to manage the participants. Two typical examples are the needs of a group to manage participants to a project and to manage the participants to a tutorial session.

This service will need to be based on strong security mechanisms, in particular with respect to authentication and authorization. During the implementation phase of SLICES, issues related to the usage of general-purpose authentication mechanisms and/or those provided by EOSC and GEANT (eduGAIN) will be investigated. Indeed, a secure but simple to use procedure to create an account is seen as an important point for user attractiveness.

This service shall as much as possible transparently implement the access policies of SLICES as all users are not expected to have the same authorization levels.

Once authenticated and authorized, users will have access to several services to manage their experiments and corresponding data.

5.5.2 Documentation and online experiment helpdesk

SLICES will be a large, diverse, and advanced experimental platform. An online documentation of the available services is an important service. A challenge will be to organize this documentation and to maintain it up to date with the evolution of the platforms, both in terms of hardware and software. This documentation will have to cover needs of beginner and advanced users.

A specific accompanying documentation for the platform will be an online helpdesk. It will be developed for assisting experimenters in the entire process of accessing the infrastructure, deploying an experiment and accessing the respective results. The helpdesk will be reachable through the Web, with pre-defined common questions and the respective actions. The helpdesk will be further enhanced through an online AI based chat bot platform such Rasa, allowing the users to interface the platform in a natural manner, and get immediate support for their experiments.

5.5.3 Accountability and Billing

This service has the task of accounting for all the usage of the platform by users with respect to SLICES policies. It covers all types of resources (hardware, software, and data). The service will provide to the user her actual usage of the platform and inform her when needed (user or system thresholds).

As some users will access the platform under a "market-driven" type, this service has also to handle the current invoice. Various modes of billing could be envisioned such as pre-paid or not. As SLICES will last many years this service has to be flexible enough to adapt to inevitable policy evolutions and special situations.

5.6 Resource management services

5.6.1 Resource Discovery and Description

A key service of the platform is the ability to discover and describe the resources it manages. The resource discovery part will be exposed in several flavors such as resource listing, browsing, or querying. As SLICES will be a large platform with many diverse components (computing unit, networking unit, IoT, storage, data, etc.), a critical property of this service is to remain user friendly, accurate, and efficient. During the implementation phase, an analysis of best practices, existing APIs, and tools will be conducted.

The second part of this service is the description of the resources. Not only, this service shall accurately describe the resources and how to use them, but it has also to manage this information over time. Indeed, during the lifetime of the platform, new hardware will be added to deal with new research opportunities. However, because of the short lifespan of hardware (3-10 years), many elements will have to be replaced. It is important to maintain a very accurate description of the platform over time for reproducibility challenges, through automated tools for resource discovery and replacement.

According to the FAIR principles discussed in Deliverable D4.2 [D4.2], a particular case of resources is the data. This service will also have to be able to query and list data that have been imported into SLICES or generated by a previous experiment. It will imply authorization issues as some data may have access restriction. Hence, it will need a tight integration with the Data Management Service described in Section 5.8.

5.6.2 Resources Reservation

At a basic level, a user will be able to reserve some resources to run her experiments on them. The amount and types of resources she can reserve will depend on her usage and the platform policy. Different modes of reservation are envisioned: immediate or advanced reservations, interactive or batch-oriented reservations, exclusive or shared access to the resources, short or long-term durations, etc. Such a variety of reservations modes have shown to be needed to cover distinct user needs and to maximize the utilization of a platform.

Hidden to the users, the SLICES scheduler will accept or not user reservation requests. We plan to take care to let the system clearly explain to the user why her request is rejected (not enough resources available, not compliant to SLICES policies, etc.).

Resource reservation should be possible at different level of granularity, either by requesting specific resources or requesting any resources matching certain characteristics. The best methods and interfaces to propose this multi-level resources reservation setup will be evaluated during SLICES' implementation phase.

Additionally, to help novice users' setup their experiments, domain-specific template-based resource reservation will be a solution to ease this step. Examples of templates could be cloud, IoT or 5/6G specific reservation. Hence, it would require that users (or communities) to define and share resource reservation template requests, potentially as a specific type of data.

5.6.3 Resources Configuration

Most resources of SLICES will be configurable. It ranges from bare-metal configuration for servers, network elements, and IoT nodes, to isolation or emulation. It is a critical point of an experimental platform such as SLICES as it enables experimenters to configure the resources as they need for their experiments. Note that resources are expected to have a default configuration that targets a common case.

Configurations ranged from simple operation such as setting the value of a parameter to complex cases where an orchestrator is needed. An example of complex configuration is the configuration of several servers interconnected with a virtual network spanning several SLICES sites. The network could include network emulation features to simulate performance degradation; the servers are provisioned using system images picked from a collection able to meet various user requirements.

5.6.4 Resource Monitoring and Profiling

In order to support large-scale experimental research, SLICES will provide an end-to-end monitoring and profiling service. This service will support the collection of monitoring and profiling data and will propose a framework to store and expose such data lively or in a post-mortem mode. This service will support several performance metrics for hardware and software components: usage of resources (CPU, memory, network, etc.), Quality of Service (latency, budget, etc.), energy consumption, thermal dissipation, etc. Monitoring is not the only target as this service will also provide some profiling components allowing to observe and analyze lively the behavior of SLICES use cases and applications over the time while managing environment parameters.

This service will handle several challenges like the interfaces with a heterogeneity of observed resources. The implementation of SLICES will design and implement the needed APIs and frameworks to support this service. It could be based on already implemented components like Kwollect [KW21] or OMF (cOntrol and Management Framework [OMF10]). These frameworks will be orchestrated at large scale by SLICES in order to support large-scale data collection.

5.7 Data oriented services

5.7.1 Data Management Service

SLICES will provide a common and interoperable service for managing data. It will be operated under common Data Governance and Management Policy (DGMP) and complying with the common Data Quality Assurance Framework (DQAF) as further detailed in Deliverable D4.2 [D4.2].

This service will manage the information/data model that defines the specification of data types and formats that SLICES generates and/or collects, and relations between data and related processes. It will provide functionalities and tools to manage the data publication and sharing, internal and open access, the data preservation and maintenance processes, quality assurance process, support data provenance, etc.

The definition of the data/information model will provide a foundation for the consistent data and metadata management and interoperability. The data model will include schemas and formats for data collected from SLICES, experimental equipment and applications. In order to ensure an appropriate quality management of data collection and sharing, the data quality assurance process will be applied to test the internal validity of the data collected both during and after the data is entered, and by making sure that proper documentation is available to the analysts who will be using the data.

Data Governance and Management strongly depends on organizational models and proper definition of organizational roles responsible for data management. A special role of the Data Manager or Data Steward will be established at each partner organization on the full time, part-time or shared basis. They will be in charge of implementing DGMP and Data Management Plan including the data quality management and data documentation. This role will also be responsible for checking that provided and collected data conform the proper national and international regulations, do not infringe IPR, and do not disclose sensitive information for publishing. The data governance and management are a shared responsibility between partners and aims to provide a robust and trustable solution for the collection, processing, and publication of quality data by exploiting the local competencies of each partner.

The Data Management Service need to be technically and semantically interoperable with the EOSC federated data infrastructure and support corresponding data and metadata formats, protocols and APIs. This service will include and rely on four underlying services: data storage, data transfer, data persistent ID generator, and metadata generator, which are described hereafter.

5.7.1.1 Data Storage

Experiments usually need input data and produce results, i.e., output data. As in other existing platforms, different types of storage systems will be provided such as a limited but rather stable storage and a larger but more volatile storage, typically to handle a large amount of temporary data produced during an experiment.

As SLICES will be a distributed platform, the storage will also be distributed. Hence, a user will be able to precisely select where input and output data for an experiment are located as distinct locations may result in distinct latency and/or bandwidth and this may impact the result of the experiment. However, it is expected that the service will ease a lot the management of the movement of data, such as providing a location-transparent data access whether it is not important for an experiment.

Several systems already exist with different tradeoffs. The implementation phase will investigate them, in particular from the point of view of privacy and performance as SLICES has to be able to deal with big data experiments.

SLICES is not expected to provide long-term storage for user generated data. Hence, user will need to move their data out of SLICES when experiments are done, typically to a user or community data storage place that will be typically part of EOSC.

5.7.1.2 Data Transfer

Users need to transfer input data to the platform and output data such as results from the platform. This traditional service will need to provide efficient and safe mechanisms to cope with the potential huge amount of data to transfer and to be able to deal with data provenance issues. Moreover, interoperability is an important aspect to ease the use of the platform. A point of particular attention will be the connection with EOSC for example as SLICES is not intended to provide long term storage outside the needs to conduct experiments. Of course, this service will have to support the implementation of FAIR principles.

5.7.1.3 Data Persistent ID Generator

A central element of the FAIR principles of EOSC is the ability to identify a piece of data that need to be presented as a FAIR Digital Object (FDO) which is identified via Persistent Identifier (PID) (see Deliverable D4.2 [D4.2] for details about DO and PID). During an experimentation many data will be generated. Hence, in compliance to best practices that will be set up in EOSC, SLICES will provide a service to generate the PID of the dataset (or FDO) as a result of an experimentation. The precise definition of this service is left to the SLICES implementation phase as it depends on advances within EOSC with respect to PID services definition. As mentioned in the resource discovery service and data management service, SLICES will provide simple mechanisms to let user register generated data that can be shared and re-used inside and outside SLICES community.

5.7.1.4 MetaData Generator

Metadata are essential to enable reuse, facilitate interoperability with external platforms and services, such as EOSC, and maximize impact. Consistent metadata specification is a key element of the FAIR Digital Object definition and supporting PID infrastructure to be built by EOSC. One of the main objectives of SLICES is to facilitate collaboration of researchers/practitioners, both within SLICES and beyond, by developing an efficient and interoperable metadata repository. A dedicated resource discovery component will allow for searching data resources through appropriately defined metadata profiles.

Metadata generation will be provided by two main mechanisms: (a) user-defined; and (b) automatically generated. User-defined metadata generation will be provided through suitable mechanisms and APIs that will allow users to enter compulsory (i.e., minimum metadata for ensuring that the object is FAIR and specific metadata related to its type, such as publication, experimental data, service or software) or optional (i.e., discretionary or recommended metadata, which can be used to improve the object's discovery and interoperability). Automatically generated metadata will be facilitated by metadata extraction components that will employ indexing and information retrieval techniques to complete the metadata generation process. Where appropriate machine learning techniques will be employed to assist the process.

The resource discovery component will utilize the above-mentioned generated metadata to support discovery and access allowing users to use complex queries, such as free-text search, combine filters and drill down to the individual components of each metadata property, to retrieve the desired digital objects.

5.7.2 Experiment data validation and correlation with other experiments

In the context of this service, the datasets produced by an experiment in a given test site could be optionally correlated with other datasets generated by similar experiments in various configurations and stored in the SLICES data storage service for instance. This feature will ensure a fast analysis of the produced experimental results and their immediate validation. Typically, the measurements of previous experiments can be retrieved and compared with the new ones realized by the experimenter who can examine in more details the highlighted differences with the previous results. Therefore, the experimenter can concentrate himself to the unexpected measured values and investigate the root cause of the abnormal values. Such correlation of results for same/similar experimental scenarios can demonstrate the differences between emulation vs large scale real-world experimentation, and prove in practice the value of SLICES-RI.

The experimenter will be able to visualize the data generated by her experiment through a tool like Grafana and then, to immediately publish openly the results in order to boost the research process. Such option will be optional, and provided as a service, along with the usual experiment visualization tools. Different means of open data access will be provided such as an Open Data CKAN server managed by SLICES or EOSC at European level. Of course, this service will employ the data transfer service, in particular for the data to be stored during a long period of time in EOSC.

5.8 Experiment management services

5.8.1 Experiment management

A user can be involved in several experiments, an experiment gathering all the needed artifacts (data, workflows, notes, etc.). This service provides an interface to manage (create/destroy/organize), browse, explore her experiments, but also experiments from other users given she is authorized. The service also deals with the management of metadata associated to the experiments to be later defined such as a description, access right, associated data, etc.

5.8.2 Experiment Control and Orchestration

A user has a full control over an experimentation; automatization can be achieved with the API of the services. However, it will require an external machine to execute this automatization. That is why SLICES will also provide three distinct types of service to control and orchestrate experimentation. A first type of experiment control and orchestration service will be to provide notebooks (like Jupyter for example) where experimentation orchestration could be defined and executed. This approach is widely deployed and many communities are using it.

A second type of experiment control and orchestration service will be higher-level experiment design and execution, such as for example to describe a campaign, made of many experiments and that may last for a long period. This type of service could be exposed through the web portal, API, or notebooks. Its particularity is that it has to manage a state and activity over a long period of time and hence it has to be resilient to platform failure and unavailability.

A third type of experiment control and orchestration service is domain specific experiment description. Sharing low-level notebook scripts is always possible, but it may be more efficient (at least for some classes of experimenters) to share configurations/artifacts at a higher level to ease experiment replicability across several platforms.

The support of domain specific experiment control and orchestration will facilitate the access, the creation, and the control of experiments, specifically to the novice users of the low-level infrastructure programming. Such software will permit the description of the use cases and scenarios without diving into complex programming and debugging. Therefore, SLICES will provide libraries and frameworks to enable the development of such automatic code generation for experiment orchestration and collection of results. The automatic code generation will use Artificial Intelligence mechanisms and Continuous Integration/Continuous Development (CI/CD) as supporting technologies. CI/CD allows the immediate validation of the produced code. Experiment requirements will be gathered through for example YAML/JSON descriptions, questionnaire forms from the users, or through a smart assistance bot such as Rasa to output the respective code/packages that need to be deployed in the platform and/or to select from the available repositories for functions/node images as well as advanced services. It will consolidate state-of-the-practice tools (such as monitoring stacks, network traffic shaping, load injector, etc.). Subsequently the user will only have to on-board the respective functions to the SLICES infrastructure in order for the experiment to begin. Similar methods will be deployed for accessing the results from each experiment run.

5.8.3 Dashboard

The dashboard service will give a complete view of the current situation of the platform as well as of previous and ongoing experiments. It will have to integrate most of the previous services for example to move data, start or cancel an experiment, post-process data, etc. It will have to provide a notification mechanism to inform users from expected events (such as the end of an experiment) or unexpected events (a fatal error occurs during an experimentation).

6 Examples of usage of SLICES Services

The previous section has described independently each service of the initial list. The goal of this section is to highlight their usage in four typical basic workflows: a basic experiment, an orchestrated experiment, a data-intensive experiment, and a course.

6.1 Basic Experiment

After connecting to SLICES, a basic low-level experiment consists in discovering the adequate resources, booking and configure them, before launching the experiment script and downloading the result before releasing the resources. Note that in this case, a direct access to the resources is assumed to launch the experiment script and retrieve the result. Figure 2 illustrates it.

This basic experiment focuses on almost the minimal services to execute an experiment. Of course, other services like experiment management, documentation, dashboard are likely to be use.

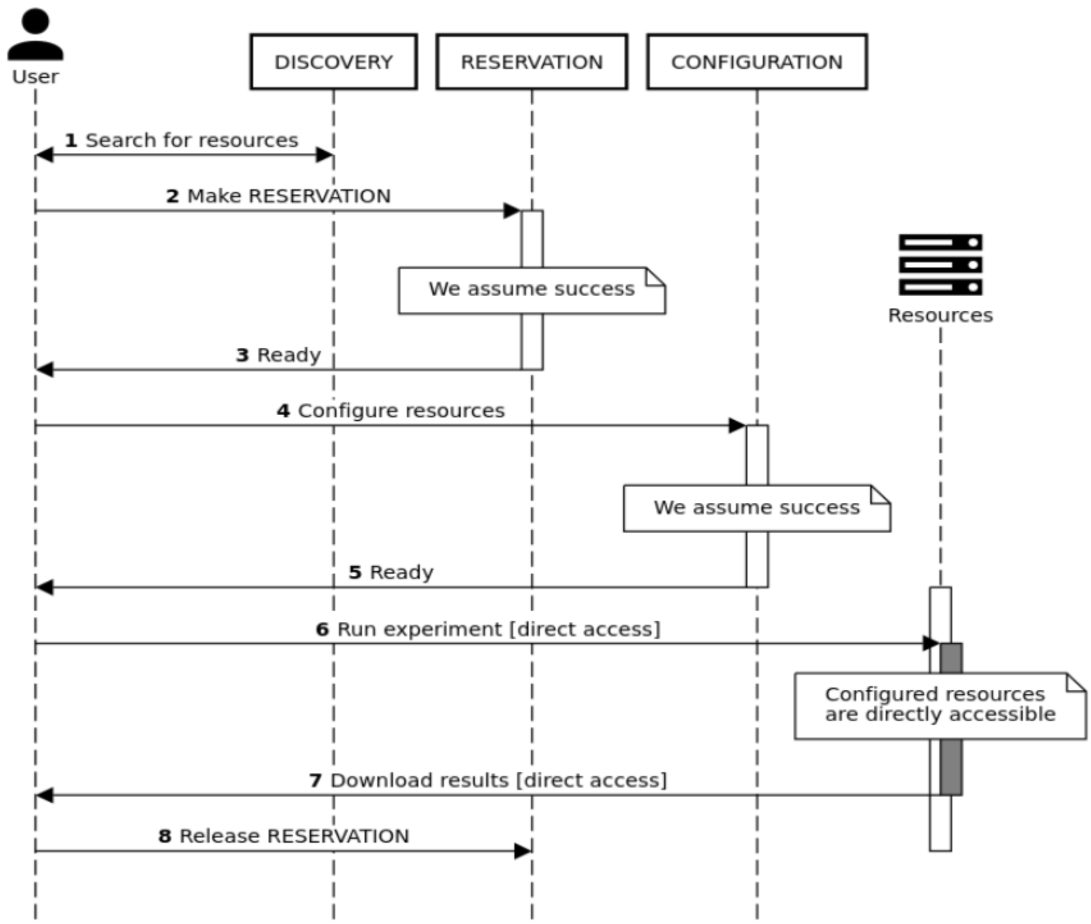

Figure 2: Basic experiment

6.2 Orchestrated experiment

A variation of the previous basic experiment is the usage of the orchestrator to execute an experiment with several configuration of the resources as illustrated in Figure 3. In that experiment, resources are potentially reconfigured during the execution loop.

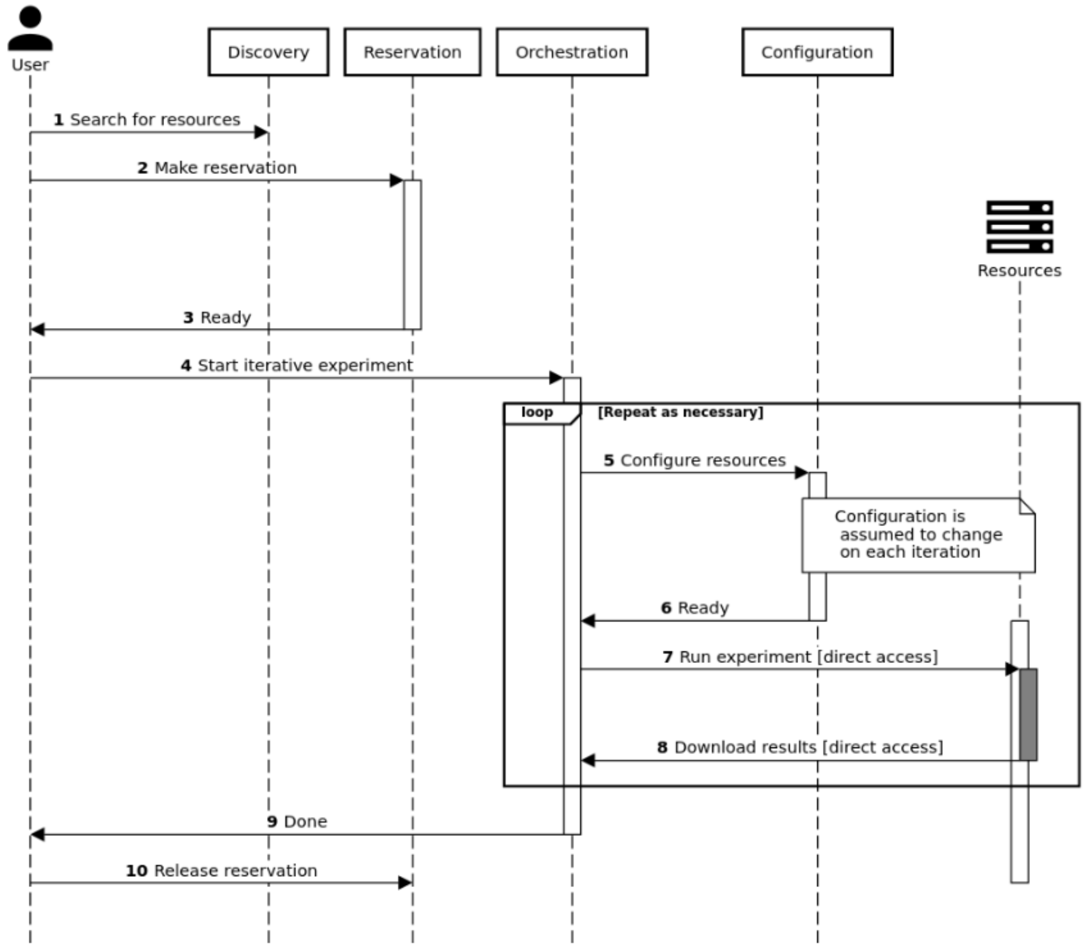

Figure 3: Orchestrated experiment

6.3 Basic data intensive experiment

Figure 4 describes an experiment where a user imports some data into SLICES before launching her experiment and then exporting the output data. The sub-services of data management, i.e., data transfer, data storage and data persistent Id generator have been displayed separately but in the same Data Service group to better illustrate them.

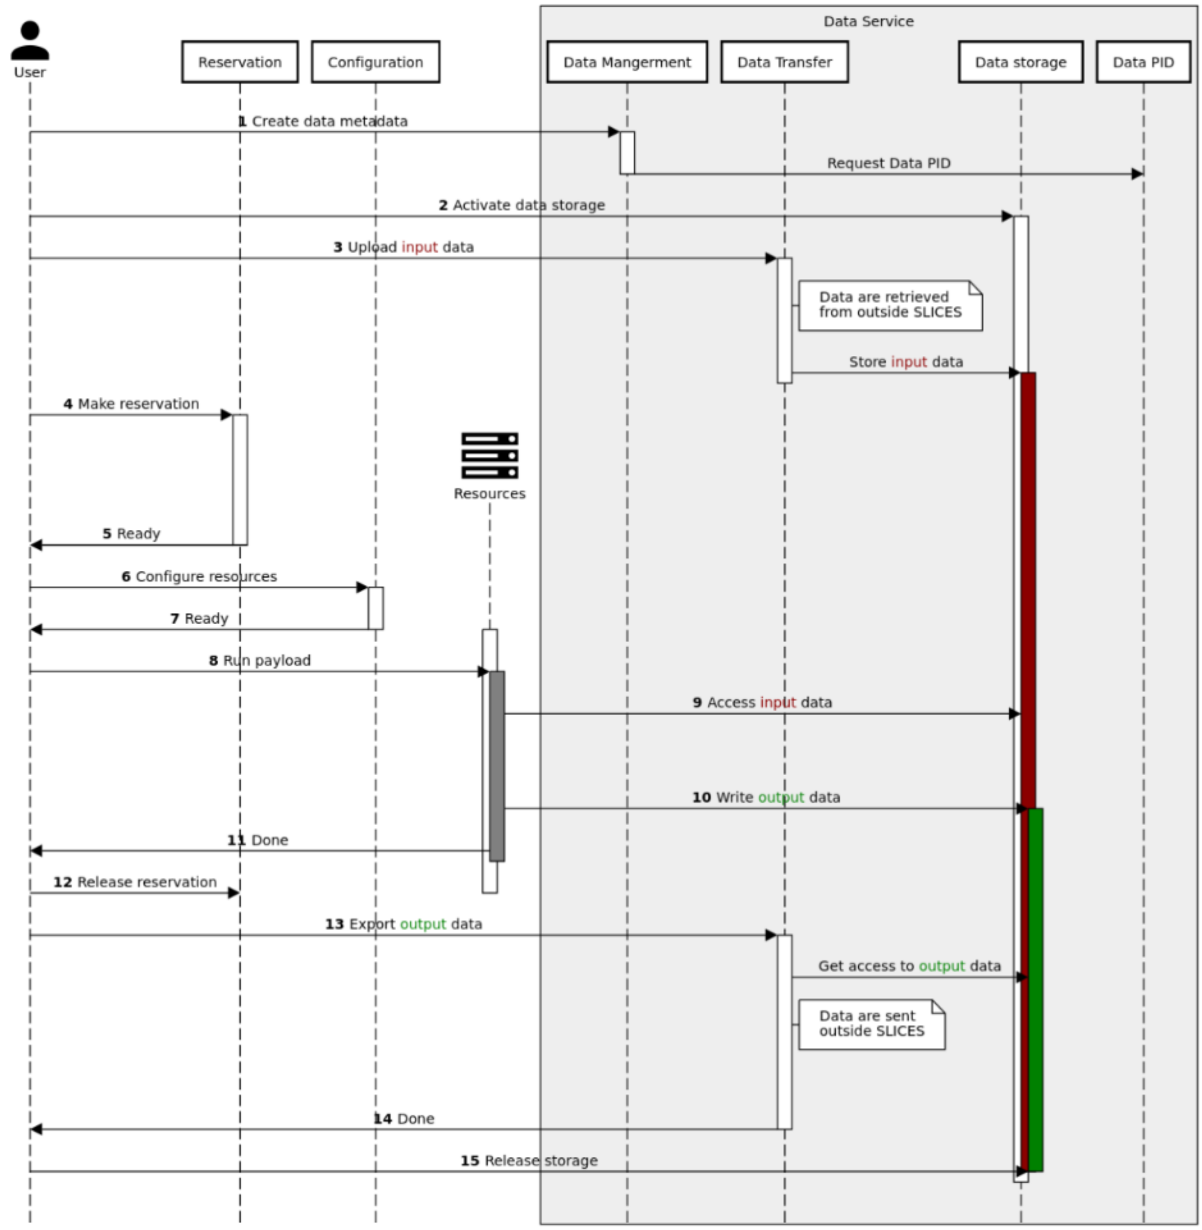

Figure 4: Data intensive experiment

6.4 Basic course

Figure 5 illustrates a typical course. First, the teacher as to prepare the course and later the participants can access the prepared environment during the actual session. Therefore, the teacher can make use of the experiment management service to manage all artifact elements needed by the course. She also makes use of the user management service to create a special group and register attendants to it. Hence, attendants would be able to connect to this particular experiment.

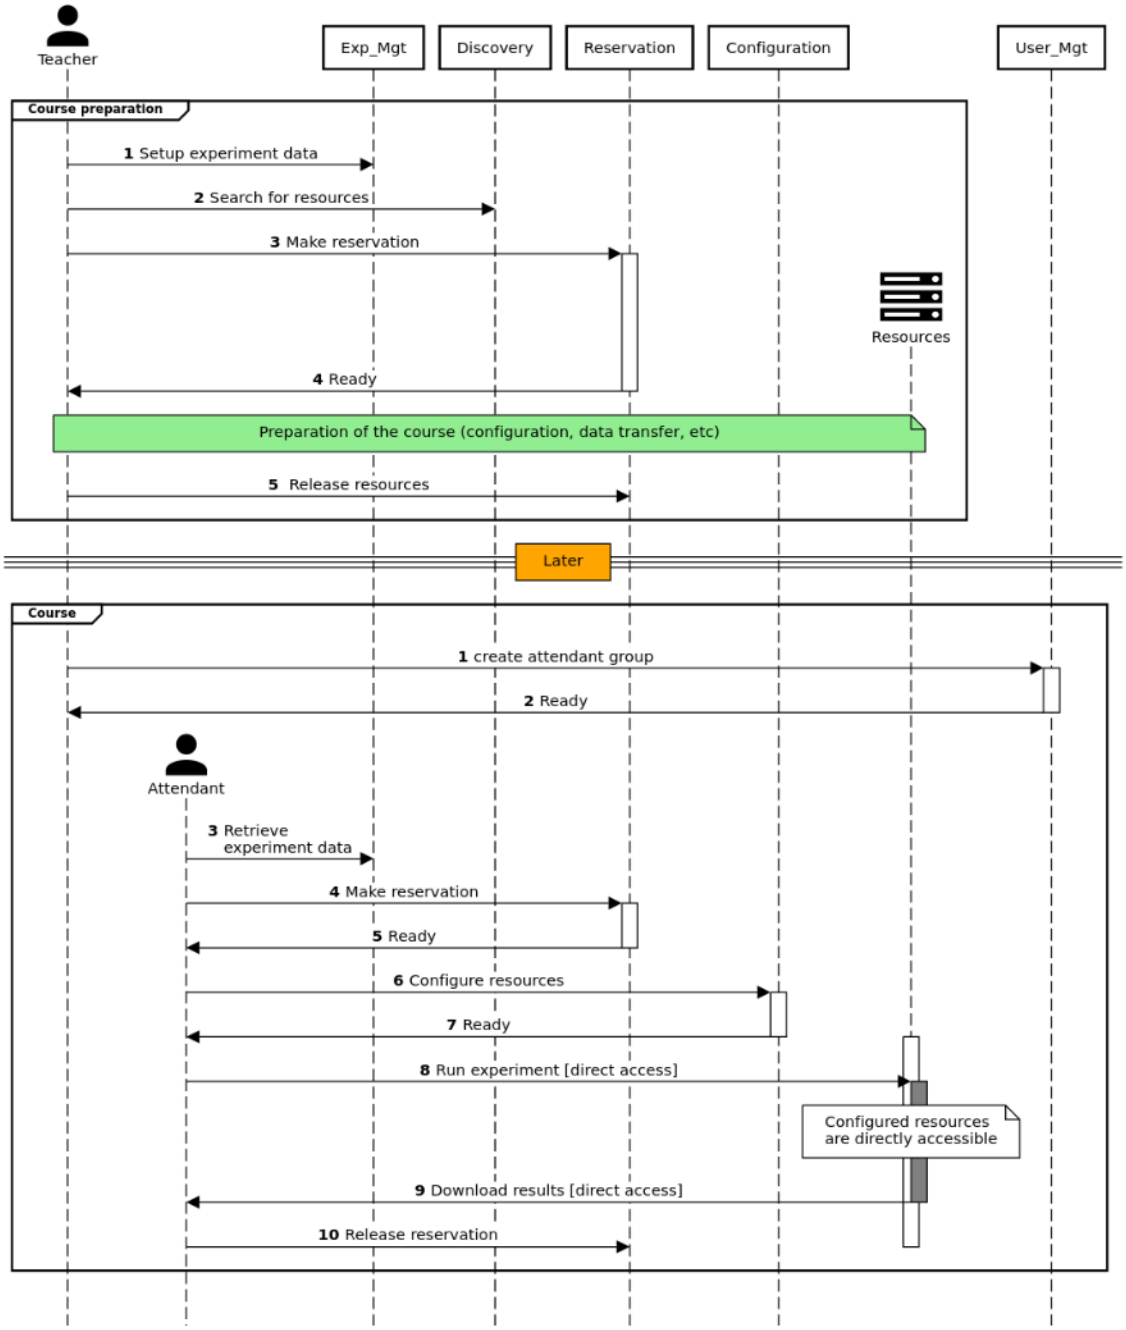

Figure 5: Basic course

7 Conclusion

The present document has focused on the services of the SLICES RI, and the manner in which SLICES will provide its experimentation offering to the relevant communities. After an analysis of user requirements and usages, the document focuses on SLICES services. First, it presents the modes and types of access that have been designed in accordance with the ESFRI guidelines, as well as with the various interfaces to the platform. Then, it deals with how SLICES will be made available through an initial set of twelve services that will facilitate the ease of access and the control of experiments, regardless of the complexity of the underlying equipment. The services are based on the long experience of the participating members in managing such a large-scale infrastructure.

8 Bibliography

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

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

[KA13] Emmanuel Jeanvoine, Luc Sarzyniec, Lucas Nussbaum. Kadeploy3: Efficient and Scalable Operating System Provisioning. USENIX; login:, volume 38, number 1, pages 38-44 (February 2013).

[GK13] Daniel Balouek, Alexandra Carpen Amarie, Ghislain Charrier, Frédéric Desprez, Emmanuel Jeannot, Emmanuel Jeanvoine, Adrien Lèbre, David Margery, Nicolas Niclausse, Lucas Nussbaum, Olivier Richard, Christian Perez, Flavien Quesnel, Cyril Rohr, and Luc Sarzyniec. Adding virtualization capabilities to the Grid'5000 testbed. In Cloud Computing and Services Science, volume 367 of Communications in Computer and Information Science, pages 3-20. Springer International Publishing, 2013.

[CL16] D. Duplyakin and R. Ricci. Introducing configuration management capabilities into CloudLab experiments. 2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), San Francisco, CA, 2016, pp. 39-44, doi: 10.1109/INFCOMW.2016.7562042.

[KW21] Simon Delamare, Lucas Nussbaum. Kwollect: Metrics Collection for Experiments at Scale. CNERT 2021 - Workshop on Computer and Networking Experimental Research using Testbeds, May 2021, Virtual, United States. pp.1-6.

[OMF10] T. Rakotoarivelo, M. Ott, G. Jourjon, and I. Seskar. 2010. OMF: a control and management framework for networking testbeds. SIGOPS Oper. Syst. Rev. 43, 4 (January 2010), 54–59.

[JUP] Jupyter notebooks https://jupyter.org/

[D1.1] SLICESA-DS, D1.1 "Technological status and capabilities of existing ICT Research Infrastructures", March 2021.

[D1.2] SLICES-DS, D1.2 "Requirements and needs of scientific communities from ICT-based Research Infrastructures", Aug. 2021.

[D2.1] SLICES-DS, D2.1 "Initial description of the SLICES architecture ", 2021.

[D4.2] SLICES-DS, D4.2 "SLICES infrastructure and services integration with EOSC and Open Science (initial proposal) ", Aug. 2021.

slicesDS