# SLICES-RI Pre-Operation Methodology and Services


1st Serge Fdida*
*LIP6 laboratory*
*Sorbonne Université, CNRS*
Paris, France
ORCID 0000-0002-6993-101X

2nd Panayiotis Andreou
*InSPIRE Research Centre*
*UCLan Cyprus*
Larnaka, Cyprus
pgandreou@uclan.ac.uk

3rd Nikos Makris
*Dept. of Electrical and Computer Engineering*
*University of Thessaly*
Volos, Greece
nimakris@uth.gr

4th Damien Saucez
*Centre Inria d'Université Côte d'Azur*
Sophia-Antipolis, France
ORCID 0000-0002-1137-6236

4th Sebastian Gallenmüller
*TUM School of CIT*
Technical University of Munich, Germany
ORCID 0000-0002-9024-0117

4th Brecht Vermeulen
*Department of Information Technology*
Ghent University/imec, Gent, Belgium
ORCID 0000-0001-7099-8687



***Abstract*—SLICES-RI is being developed as a scientific instrument, following established methodologies and best practices from the scientific community. This approach is increasingly essential to address the accelerating pace of data-driven research, which continuously generates a deluge of publications. The core mission of a scientific instrument is to provide a standardized and trusted reference, enabling direct performance comparisons of algorithms and ensuring reproducibility — thereby simplifying and strengthening the peer review process. To meet these goals, SLICES-RI adopts an intent-based design and generates tailored Blueprints for specific scientific questions. While it does not aim to be exhaustive, this approach significantly reduces the complexity involved in designing and executing experiments. In this presentation, we share the current deployment status of SLICES-RI as it entered its pre-operational phase. We illustrate its capabilities through the Post-5G and Federated Learning Blueprints. Additionally, we highlight MRS/DMI, a key feature of SLICES-RI, which offers an advanced data management framework aligned with the FAIR principles and fully integrated into the experimental workflow.**



***Index Terms*—Experimentally-driven research, Research Infrastructure, FAIR data management, 6G, Cloud.**


## I. Introduction

SLICES-RI [1] is a European scientific instrument developed under the umbrella of the European Strategy Forum on Research Infrastructures (ESFRI). It entered the ESFRI roadmap in 2021 and is currently in its preparation phase. As any other ESFRI project, SLICES-RI is defined and designed as a scientific instrument. Those are of utmost importance as a large fraction of scientific papers published in the past decade involved the use of advanced scientific instruments. The value of adopting common tools for a research community exploring some specific scientific questions is that it provides a reference environment for the study, analysis, and testing of various assumptions and solutions. As such, it accelerates the pace of scientific progress and depth of understanding because of this common ground, helping scientists to measure the subject of their investigation. In addition, a scientific instrument comes with a methodology and practices that the community should adopt and follow.

The use of well-identified scientific instruments by research communities varies from one domain to another, based on their practices, experience, and maturity. This diversity of practices also exists within the computer science and electrical engineering community, some being more advanced such as the famous Lena (now banned for good reasons), widely adopted by the image processing community to test various algorithms. Its value lies for long in the fact that it became a common reference image adopted by the early computer vision and image processing community. As such, it became a de facto standard, considered as a reference image enabling direct performance comparison of algorithms and reproducibility, making peer review simpler. However, such an agreed reference does not exist in many fields and in particular in the domains covered by SLICES-RI, namely digital infrastructures, including networking, wireless, cloud, and IoT. This weakness is becoming even more apparent due to the stiff increase in the number of papers published, in particular when using AI/ML. This concern echoes to the seminal paper of J. P. A. Ioannidis published in 2005 entitled "Why most published research findings are false" [2]. The paper even conjectures that the broader a research domain, the higher the likelihood for a result to be false. To complement this concern, reproducibility in computer system research is making very little progress. As outlined by C. Colleberg and T. A. Proebsting [3], one way to circumvent this severe concern is to reward commitments to sharing research artifacts. This is also a mandatory service that a research instrument should


This work is partly funded by the European Commission (Horizon Europe programme) under the Grant Agreement numbers 101079774 (SLICES-PP) and 101139257 (6G-SUNRISE). Asterisk indicates corresponding authors.

provide.

## II. Background and Related Work

Internationally, the need for advanced and sustainable test platforms to research in digital infrastructures has been recognized for years. Initiatives such as the NSF's GENI project ($120M, 2008–2016) and the European FIRE program (200M€, 2007–2022), structured efforts to build research infrastructures in this domain. These were followed by further initiatives such as NSF PAWR, Chameleon, and Fabric in the US, CENI in China, and the Stream C program under Europe's SNS-JU.

Despite these significant achievements, a recurring challenge remains: the long-term sustainability of such platforms. Limited operational lifespans often constrain their ability to deliver continuous value, ultimately diluting the return on the substantial human and technical efforts invested in their deployment and maintenance. Addressing this challenge is critical to ensure that future infrastructures can fully support the dynamic and evolving needs of network and system research.

## III. SLICES-RI Pre-Operation

A key strength of SLICES-RI lies in its long-term sustainability and its vibrant, pan-European research community. Currently, SLICES-RI is backed by 16 European member states, with formal political support secured from 12 of them. This strong foundation reflects a collaborative commitment between the European Commission and national governments, adopting the model applied in large-scale scientific instruments such as telescopes or colliders — infrastructures that typically require years of coordinated effort from design to full operation and $B investment.

In addition to providing virtualized resources and services such as compute, storage, and connectivity, SLICES-RI also aims to address the scientific value embedded in the research data and digital artifacts it produces. Managing this wealth of information in accordance with European Open Science policies and FAIR principles (Findable, Accessible, Interoperable, and Reusable) is a core objective. This requires the implementation of well-defined metadata standards, including experiment descriptions, to ensure that the data remains valuable, shareable, and reproducible across the scientific community. It also provides a key support for reproducibility. As a consequence, SLICES-RI has designed a Metadata Registry System (MRS) and a Data Management Infrastructure (DMI) to enforce the FAIR principles [4].

SLICES-RI will offer a broad and diverse range of resources, most of which will be provided and managed by its national nodes. The types of resources mobilized will reflect both the shared objectives outlined in the SLICES-RI roadmap and specific national research interests, creating strong synergies between them.

Now that it has entered its pre-operational phase — a stage designed to prepare for full operational status following ESFRI guidelines — SLICES-RI has already deployed a first set of resources, which will progressively grow and evolve over time. This portfolio will include computing, storage, and communication components, ranging from network infrastructure and wireless systems to IoT devices and bare-metal environments.

Given the diversity and complexity of these resources, it is neither feasible nor desirable to expect experimenters to manually manage and configure them for each experiment. Describing, selecting, and exposing resources in a way that matches both the experiment's objectives and the user's expertise is a longstanding research challenge. This is due to a diversity of needs from very detailed research questions on a specific component to end-to-end analysis of vertical application domains that render the identification of the scope of resources of utmost importance.

To address this, SLICES-RI adopts an intent-based approach: the system tailors the set of accessible resources based on the scientific question, the experimental methodology, and the target research community. To define this approach, the project engages directly with multiple research communities, identifying their most pressing challenges and translating them into reference environments that enable meaningful testing and direct performance evaluation. This process delivers *Blueprints* for that specific community.

## IV. The Blueprint Concept

Many domains within Computer Science and Electrical Engineering still lack widely adopted and standardized reference platforms for conducting experiments. The Blueprint concept aims to address this gap. Blueprints extend beyond individual experimental setups by providing both resources and structured workflows, collectively forming a reusable platform for executing scientific experiments.

By leveraging Blueprints, researchers rely on a functional and validated experimental environment, eliminating the need to build and configure customized testbeds prior to pursuing their scientific objectives. This enables researchers to concentrate their efforts on addressing core scientific questions, rather than investing significant time in infrastructure preparation.

To be truly effective, Blueprints must address two fundamentally distinct requirements: (1) *community-specific needs* and (2) *universal scientific practices*.

Community-specific needs refer to domain-dependent requirements such as hardware configurations (compute, storage, memory, network), software environments (bare-metal, virtualized, or containerized platforms), and compliance with regulatory constraints (e.g., licensing, radio spectrum usage). Blueprints must be tailored to reflect these requirements for each scientific community.

In parallel, Blueprints must support universal scientific practices, which are applicable across all research domains. These practices include the reproducibility of experiments, as well as the systematic creation, management, and sharing of experimental data and metadata. Given their generality, solutions addressing universal practices must be designed for broad applicability across various Blueprint implementations.

Engaging with the relevant research communities, SLICES-RI produced three Blueprints: Post-5G, Federated Learning, and Edge-Cloud Continuum. The first two are now in production and can be tested and discussed. More Blueprints are currently under investigation as we succeed in mobilizing the related research communities.

## V. Post-5G Blueprint experimental research

The SLICES-RI *Post-5G Blueprint* has been designed in a flexible manner, towards providing the community with a set of tools and software that can be used to replicate and reproduce experiments. The framework translates experimentation intents (as captured by the end users) into low-level configurations for facilitating the deployment of complex cutting-edge environments. The Blueprint targets providing resources, tools, and code for beyond state-of-the-art experimentation with emerging telecommunication networks and accompanying infrastructure. To this aim, the provided resources are fully virtualized, and run in a cloud-native manner. Experimentation is supported for vertical testing with the provided networks, or delving into details and algorithms of the network, towards providing the community with novel solutions for the evolution of future networks. The accompanying workflow for experimentation has been designed in order to facilitate repeatable experimentation, with all the outputs of experiments being captured and automatically collected in the SLICES-RI DMI storage, and annotated with the respective metadata that allow easy recall, comparing, and publishing of the results. The workflow for experimenting with the Post-5G Blueprint is provided in Fig. 1 and can be summarized to the following:

1) The experimenters who want to run an experiment in SLICES-RI access a web portal, where all the key configurable parameters on the infrastructure are available. There, users can define the granularity of measurements to be collected, and the equipment (location and type) to be used for their experiment.[1]
2) Following the experiment definition, the SLICES tools and backend pack all the configurations and tools that are needed and upload them to a central server.
3) Next, the experimenters need to reserve the resources that they want to use in their experiment. The resources can be entire testbeds, or parts of them, defined by selecting the experimentation nodes during the first step. Reservations use a calendar tool, and can be scheduled for any future time that the resources are available.
4) Once the reservation is active, the experiment is uploaded to the respective nodes and executed. Execution and component orchestration is done through the *pos* controller, used to facilitate reproducibility of results across different nodes/sites [5].
5) Once the experiments are completed, the results and key monitored parameters, are collected and uploaded to the SLICES-RI DMI. Metadata annotation is taking place through the SLICES-RI MRS system, and users can access their data, postprocess them or publish them according to their demands.

The equipment available for hosting the experiments consists of various software-defined radios (e.g., USRP B210, N3xx, X410) at different locations, O-RAN compatible Remote Units (O-RUs) compatible with the 3GPP split 7.2 specifications [6] (e.g., Benetel RAN550, LITEON FR2), 5G modems (e.g., Quectel, Phones), and high-end compute infrastructure. The number of nodes with the specific equipment is constantly being updated during the pre-operations phase of SLICES-RI, adding new sites and new types of equipment, hence we suggest that the reader refers to the respective documentation[2] for exact numbers. The networks are deployed as cloud-native functions, allowing the users to interface and interact with them as in the public cloud. The infrastructure is using PTP synchronization within the sites where O-RUs are present, required for their operation. The network is using the OpenAirInterface 5G implementation for the Core Network and the RAN, and uses either emulated or real RF radios, if the users selects it.

To streamline the creation and execution of experiments with the Post-5G Blueprint, SLICES-RI provides a user-friendly dashboard (cf. Fig. 2) through which certain parameters of the experiment can be tuned. Alternatively, an LLM-based chatbot assistant is available to assist the users setup the experiment properties, by just conversing with the agent. The chatbot has been trained using the available configuration, and can produce accurate outputs, compatible with the infrastructure. Some examples from the end-user interaction with the chatbot can be observed in Fig. 3. Through these options, SLICES-RI covers users with different backgrounds.

Beginners or people without any prior experience to 5G can use the chatbot to properly setup their experiment, and experienced users can use the graphical interface or a declarative meta-model to prepare their setup and access the resources. In particular, SLICES-RI Post-5G Blueprint is also used for education and has already been experienced by students in cellular network curricula. The material for the classes and the labs will be posted soon in the SLICES-RI Academy.[3]

In order to illustrate the potential of our Post-5G Blueprint, we present two indicative examples of experiments conducted with the current SLICES-RI infrastructure. Both of them use the 5G infrastructure, and take advantage of the dynamic nature of the cloud-native functions for the 5G network to:

[1]Advanced experimenters also have access to the meta-model used behind the portal, to allow a programmatic approach for configuration.

[2]SLICES-RI documentation: https://doc.slices-ri.eu/

[3]SLICES Academy: https://www.slices-ri.eu/slices-academy/

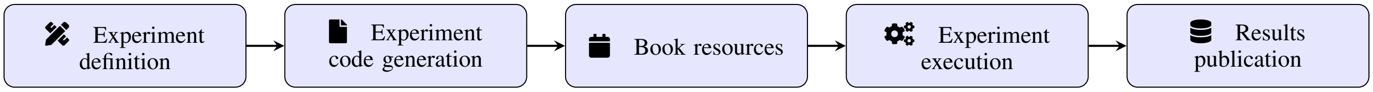


Fig. 1. Experimentation workflow with the Post-5G Blueprint in SLICES-RI

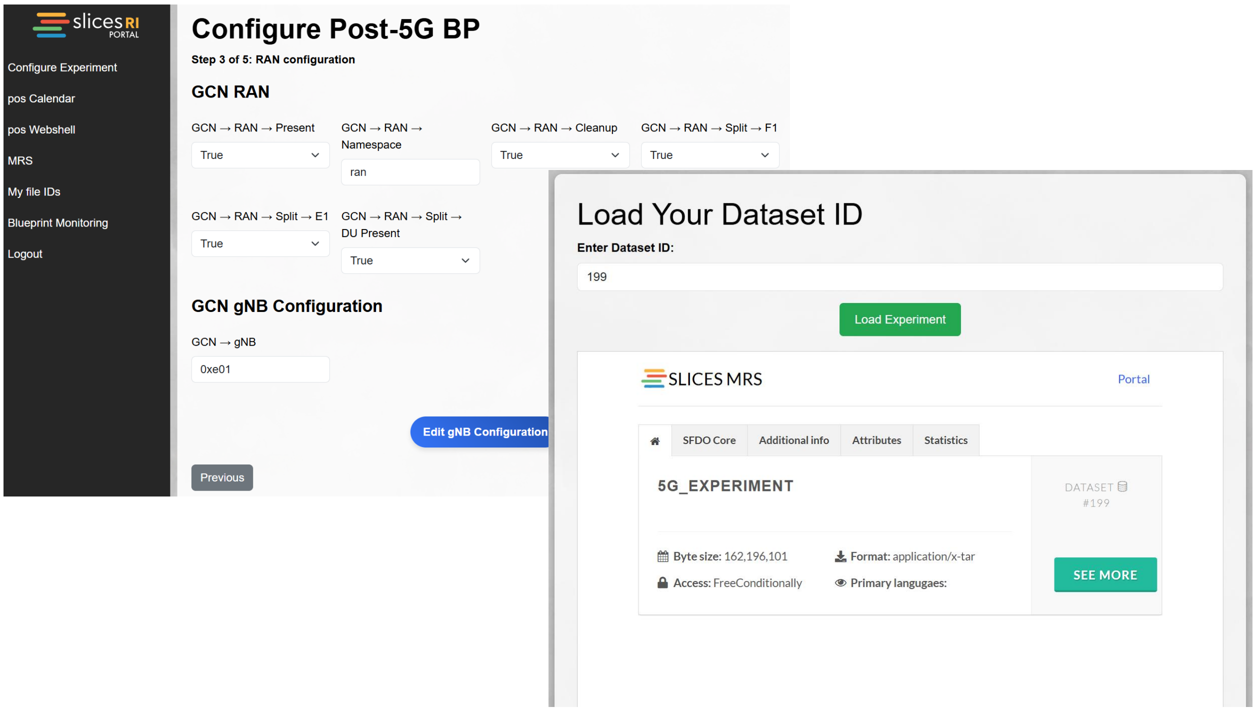


Fig. 2. Configuration dashboard for the Post-5G Blueprint, and access to the Metadata Registry System (MRS).

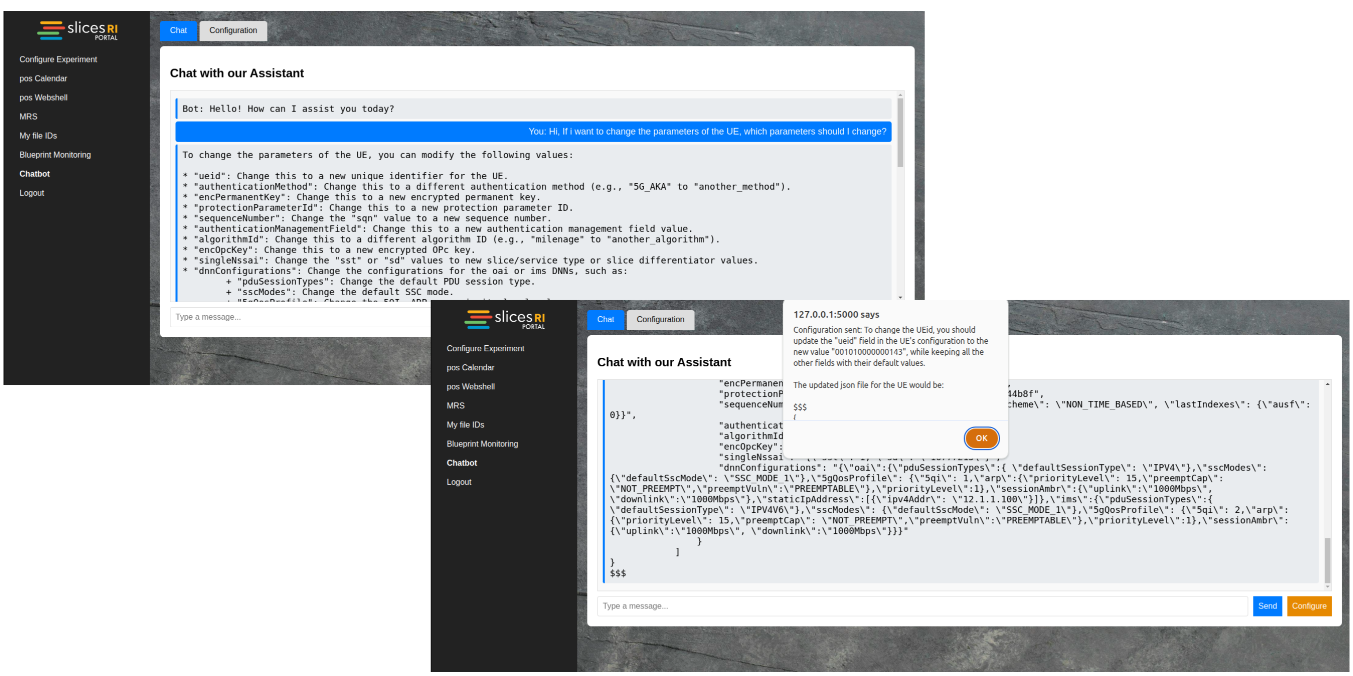


Fig. 3. Simple queries to the chatbot on what parameters can be configured; by hitting configure, the experiment configuration is sent to the backend server for conducting the experiment.

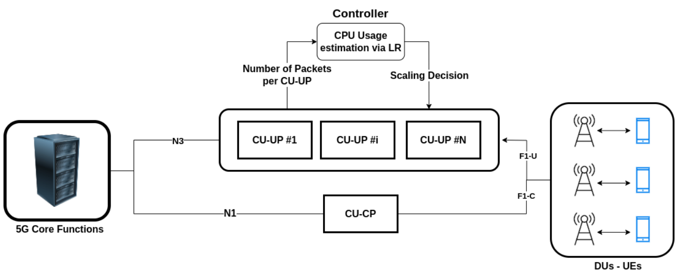


Fig. 4. Dynamic scaling of CU-UP to enhance the 5G network reliability [7].

1) Develop a dynamic scaling mechanism for the RAN, that appropriately scales the under-stress functions to accommodate user demands.
2) Dynamically switch between different User Plane Function (UPF) implementations, to minimize the overall network consumption while accommodating user demands.

### A. Dynamic scaling for the CU-UP

The high demands for network speed and data volumes that the network needs to transfer for future applications significantly stress different components of the disaggregated 5G and beyond networks. Through thorough experimentation on the SLICES-RI infrastructure [7], the CU-UP component of the disaggregated RAN has been highlighted as the data-path bottleneck under high traffic loads and the target for optimization. To this aim, a mechanism that estimates CPU usage based on a Linear Regression model has been developed, which considers the rate of incoming packets to the CU-UP instance and appropriately determines the scale of the CU-UP component to appropriately accommodate the network traffic. As the system is cloud-native, horizontal scaling is applied in a proactive manner on the CU-UP component (cf. Fig. 4), during the dynamic network operation. The work highlights the importance of adapting such a scaling mechanism to ensure the robust reliability of the network, and is taking advantage of the experimentation components available within SLICES-RI.

### B. Dynamic UPF selection

Software-Defined Networking (SDN) and Virtual Network Functions (VNFs) are key enablers of future networks, as they provide the flexibility to tailor each network based on the provider needs, and enhance it with new capabilities depending on the network operation. In this context, many implementations exist for the User Plane Function (UPF) for the 5G core network, each exhibiting different characteristics. In this context, we delve into two of the software implementations, the SPGW-U and VPP, and evaluate their performance in terms of achievable throughput and energy consumption. By integrating Machine Learning, we develop a proactive switching mechanism to dynamically switch between the two software implementations of the UPF, based on the anticipated throughput and energy consumption (cf. Fig. 5).

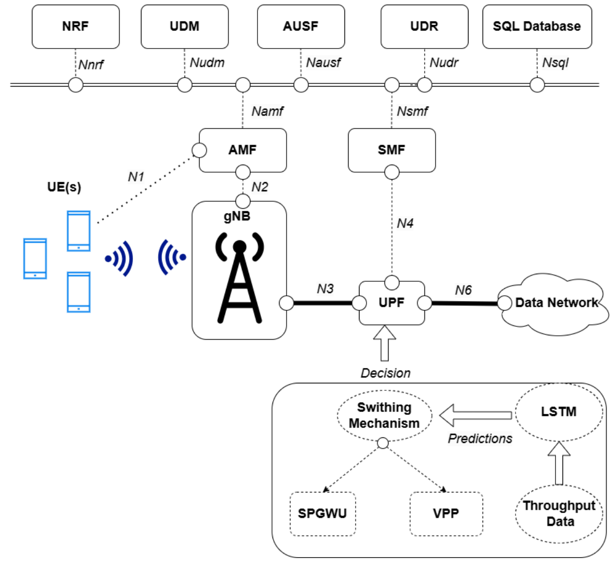


Fig. 5. Architecture to support UPF switching over SLICES-RI

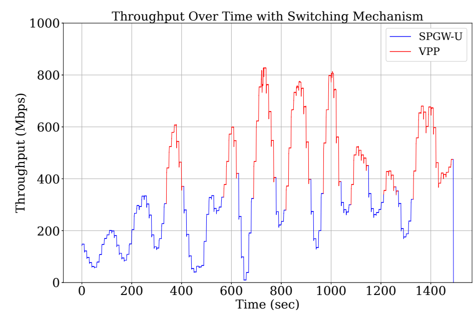


Fig. 6. Throughput enhancement when using the UPF switching scheme

SPGW-U offers limited throughput performance (restricted at about 400 Mbit/s in our setup) but consumes reasonable energy, whereas VPP delivers significantly higher throughput performance at the cost of much greater energy. The proposed mechanism runs over the SLICES-RI infrastructure, evaluates these trade-offs and employs an integrated LSTM model to proactively predict throughput values (cf. Fig. 6).

## VI. Machine Learning and Federated Learning Blueprints

In this section, we focus on the *Machine Learning* and *Federated Learning* Blueprints to illustrate the complete experimental workflow and demonstrate how researchers can leverage these Blueprints to design and execute their specific experiments. Additionally, the Machine Learning and Federated Learning Blueprint presents another significant resource offered to the community. While the Post-5G Blueprint primarily provides a programmable, hardware-based, high-performance telecom infrastructure, the core objective of the Machine Learning and Federated Learning Blueprint is to offer a versatile platform for executing machine learning workloads.

To this end, a wide range of computational resources is made available, including various types of CPUs and GPUs, all interconnected through a high-speed network, as well as high-capacity NAS and S3 storage. Efforts are made to ensure seamless integration and utilization of these resources within a unified framework.

A widely adopted framework in the machine learning research community is Jupyter Notebook [8], and therefore, the entry point for machine learning-related research within SLICES-RI is a Jupyter Hub. This environment is preconfigured with all commonly used machine learning stacks (e.g., PyTorch, TensorFlow, Julia, R, Scala, etc.) and is fully integrated with the SLICES-RI underlying hardware infrastructure. This integration allows researchers to selectively choose the appropriate components for their research. For instance, in Fig. 7, we compare the completion times of a model inference in PyTorch, both with and without the use of GPU with CUDA 12. In this example, the inference task is performed for a 10-class classification problem using a fully connected layer with 10,000 random features and a Rectified Linear Unit activation function. Randomized inference was repeated 10 times to obtain the results.

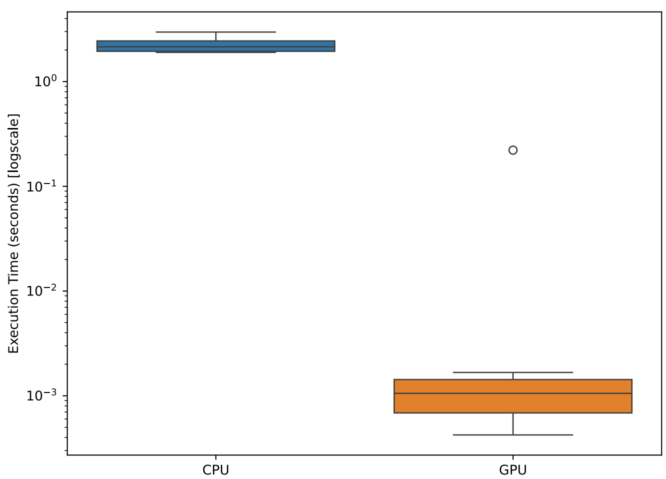


Fig. 7. Comparison of Model PyTorch Inference Completion time with or without GPU.

In contrast to other machine learning infrastructures, SLICES-RI is implemented as an open white box, offering researchers access to the underlying infrastructure and implementation. This enables them to precisely define where their workloads are executed and to configure the underlying resources according to their needs. This flexibility allows the execution of complex federated learning tasks, with components distributed across various environments, including containers, virtual machines, or even bare-metal resources. As a proof of this versatility, and as discussed in the previous section, machine learning techniques were employed to perform dynamic selection of User Plane Functions (UPFs) directly within the SLICES-RI infrastructure.

## VII. Data Management Infrastructure

SLICES-RI fully adheres to FAIR and Open Data principles enabling scientific data-sharing, in line with Europe's open science policy. To this end, all experiment results and key monitored parameters, are annotated with appropriate metadata to facilitate their findability, accessibility, interoperability, and reusability. Metadata management takes place in the Metadata Registry System (MRS) that resides within the Data Management Infrastructure (DMI).

In particular, SLICES-DMI facilitates the following:

- **Flexible Hierarchical Metadata Model (Domain-agnostic and Domain-specific)**: SLICES-RI considers the utilization of a hierarchical model consisting of compulsory metadata attributes that are domain-agnostic and can describe any digital object (e.g., data, AI models, services) ensuring conformance to FAIR and Open Science principles. Where appropriate, SLICES supports additional optional metadata attributes accompanied by their metadata model to further enhance the description of the object. The model is flexible and open, allowing extensions with new attributes or new types/categories as well as with new additional hierarchy levels.
- **Hybrid Metadata Production (human and machine-generated)**: SLICES streamlines metadata production using appropriate tools that can automatically generate/extract metadata (e.g., creation date, dataset attributes). Manual human-based metadata creation is supported (e.g., description, keywords) to enable user-defined metadata. Furthermore, predefined input validation workflows are employed to improve the quality of data.
- **Wide-ranging Interoperability**: Metadata includes attributes to support different levels of interoperability: *semantic*, which allows internal and external systems to discover and understand what the underlying object is; *legal*, which describes the restrictions in the data; and *technical*, which enables systems to communicate effectively using appropriate catalogs and services. Furthermore, SLICES-RI ensures interoperability with other systems (e.g., EOSC, aggregators) by supporting the transformation of the data to other metadata formats (e.g., DublinCore, OpenAIRE).
- **Long-term Reusability**: It is expected that through the lifetime of the SLICES project, the metadata model will evolve. The metadata model incorporates appropriate elements to describe the metadata model itself, including specific attributes, such as the metadata version, and utilize services where systems or users can obtain this info. Vocabularies and Standards utilized by specific attributes should also be versioned. Currently, SLICES-RI supports metadata crosswalks between some of the most predominant metadata standards, such as RDA metadata IG, EOSC minimum metadata set, Dublin Core, Datacite 4.3, and DCAT 2.0.

- **Enhanced Discovery**: SLICES-RI supports simple and advanced (query-based) discovery with both manual and automatic descriptors. Automatic descriptors may come from metadata data extraction (e.g., creation date, file size) but also using built-in data analysis functions (e.g., basic statistical descriptors of dataset attributes).
- **Metadata Quality Assurance**: A Metadata Quality Assurance Framework should be put in place as part of the Data Governance Group with appropriate metrics to assess the quality of metadata, its FAIRness, and other objectives.
- **Metadata Governance**: The members of the Data Governance Group should have the knowledge and authority to make decisions on how metadata are maintained, what format is being utilized, and how changes are authorized and audited.
- **Open-Source**: Currently, SLICES-RI has initiated a working group that will promote community engagement through open-source contributions, to enhance its metadata/data management and maximize its impact.

The following sections provide key insights into the underlying DMI components and their interactions.

### A. *Metadata*

Metadata is among the key enablers of open science, facilitating the repeatability and reproducibility [9] of experimental results by providing the necessary information to understand, evaluate, and replicate a study's methods, data, and findings, ensuring transparency and integrity in scientific research. Carefully crafted metadata must adhere to open science principles, such as FAIR, to facilitate human- and machine-readability and, most importantly, machine-actionability, and ensure compliance with legal and ethical regulations.

SLICES-RI implements a flexible hierarchical metadata model coined SFDO, consisting of three levels as illustrated in Figure 8. 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-RI 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 date, 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 data (e.g., wireless radio transmissions) and includes 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.

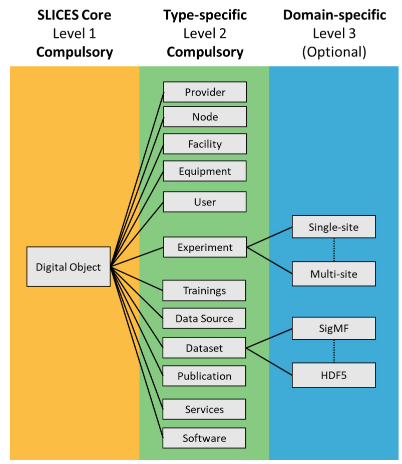


Fig. 8. SLICES Fair Digital Object (SFDO) Hierarchical Metadata Model

The SLICES-RI metadata catalog is the result of thorough analysis and crosswalks between some of the most predominant metadata standards, such as RDA metadata IG, EOSC minimum metadata set, Dublin Core, Datacite 4.3, and DCAT 2.0. Table I presents the main field categories of the metadata catalog.

### B. *Metadata Registry System*

This section presents the Metadata Registry System, which realizes the SFDO metadata model and provides access and management operations for digital objects that adhere to the SFDO structure.

*1) Architecture:* MRS provides access and management services to SFDOs using three components as illustrated in Fig. 9. First, the metadata persisted in a repository, implemented as a Postgres database. Second, a backend is responsible for exposing the repository as a REST API while providing authentication/authorization, backward compatibility, backward maintenance, and other functionality. Finally, a web portal is provided to facilitate human interaction with MRS.

*2) Repository:* MRS uses a dedicated repository for storing the metadata models and the digital objects. An important requirement for metadata persistence was to use an open-source solution. PostgreSQL was selected due to its extensive feature set, reliable track record, and widespread familiarity. However, this component can be easily replaced in the future based on new requirements, e.g., scaling, replication, etc. For example, given the nature of MRS storage and access patterns, the storage could be switched for NoSQL document and graph database engines.

TABLE I
METADATA CATEGORIES

| Code | Name | Description |
|---|---|---|
| PI | Primary Information | Important information that is used to identify and describe the resource, such as internal and external identifiers |
| MA | Management Information | Information related to the management of the SFDO, including the versioning and metadata profile used to interpret it |
| AC | Access Information | The manner and mode in which the SFDO can be accessed |
| CL | Classification Information | Classification of the SFDOin specific scientific domains |
| PB | Publication Information | Specific publication information such as the date of submission and acceptance |
| FN | Financial Information | Information related to the financial aspects of access for the SFDO |
| SU | Support Information | Support information, such as helpdesk URLs, manuals, terms of use, etc. |
| US | User Information | Information related to registered users. The information access mode is by default private and only identifiers are used |
| AT | Attribution Information | Information related to the funding instruments (e.g., body, project) of the SFDO |
| MA | Maturity Information | Fields utilized specifically with services and software and describe important information such as standards, open source technologies, and certifications |
| ST | Spatio-Temporal Information | Location and Time descriptors. Locations use Geographic coordinates when applicable |
| GE | Geographic Information | Specific geographic coordinate attributes i.e., latitude, longitude, and altitude |
| LA | Language Information | Language descriptors |
| DA | Dataset Information | Specific dataset descriptors, such as format, size |
| RT | Rights and terms of access | Legal information related to the rights of access, such as licenses, copyrights, etc. |
| SW | Software | Software-related attributes, such as repository, documentation, and programming language used to develop the software. |

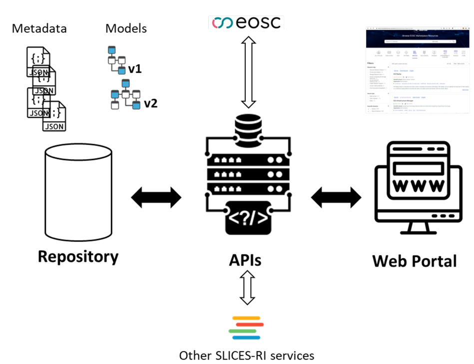


Fig. 9. Metadata Registry System (MRS) components

SFDOs are stored using a Table per Hierarchy (TPH)[4] method, i.e., SFDOs of different Level 2 types share the same database table. To accommodate for this, Level 2 attributes (columns) are set as nullable, even if they are required. As the database is only ever accessed by the backend (cf. Section VII-B3), the correct schema is still enforced before data is persisted. If the database was accessed directly by multiple systems, check constraints could be used to enforce non-nullability based on the discriminator column.

*3) Backend:* The backend is the core of the system. It is implemented using ASP.NET Core Web API, Entity Framework Core and several other widely adopted open-source libraries, such as NSwag (OpenAPI spec generator), Asp.Versioning (for API versions), QueryKit (for the advanced search), lowering the barrier for contribution. SFDO and the attributes are defined as C# classes (an abstract class for Level 1 and one class per Level 2 type, inheriting the L1 abstract class).

The backend exposes the metadata to the rest of the world in the form of a REST API. The REST API is described by an OpenAPI 3.0 specification [10], which is automatically generated from C# class and method definitions. OpenAPI (previously known as Swagger) is a widely-adopted modern standard for API description. Tools exist for consuming OpenAPI-documented REST APIs in many modern programming languages and frameworks, enabling low-barrier interaction with MRS for any potential consumer.

Besides exposing the metadata to the world, the backend provides additional functionality, such as authentication and authorization. Authentication is performed via validating Bearer JWT (JSON Web Tokens) passed as a header in the HTTP request. The JWT is issued and signed by the SLICES-RI Open ID Connect [11] STS (Secure Token Service). Using this approach, MRS, as a component of SLICES-RI, uses the same AAI as the rest of RI, ensuring a smooth user experience across the entirety of RI services.

At the moment, the authorization layer operates in a binary mode — authenticated consumers can perform mutating operations (e.g., create, update, and delete an SFDO); unauthenticated users are not allowed. As we progress in pre-operation, the SLICES-RI AAI (Authentication and Authorization Infrastructure) being developed, the MRS authorization will be inherited from the SLICES AAI.

[4]https://learn.microsoft.com/en-us/ef/core/modeling/inheritance#table-per-hierarchy-and-discriminator-configuration [Last accessed 09 February 2024]

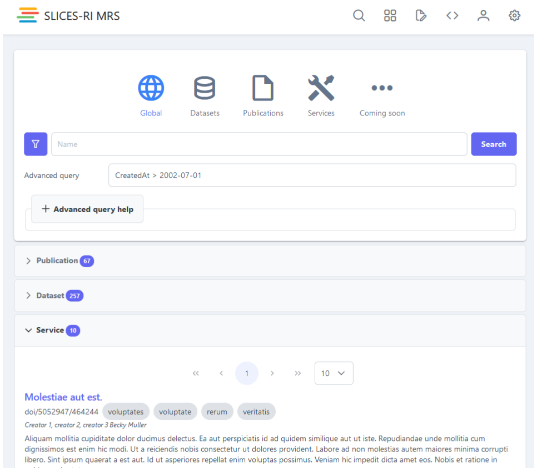


Fig. 10. MRS web portal — search functionality

All APIs exposed by the backend are versioned to provide backward compatibility. The version is specified as part of the URL path, thus ensuring that any future usage by a specific client will remain consistent. The backend handles any required translation, allowing the underlying storage format to be decoupled from any client interaction.

Any change to the metadata is accompanied by an increase in the metadata/API version number.

The API paths for SFDO CRUD (Create, Read, Update, Delete) operations are as follows:

- **Create**: POST /v[apiVersion]/digital-objects, Body: attribute values
- **Read**: GET /v[apiVersion]/digital-objects/id
- **Update**: PUT /v[apiVersion]/digital-objects/id, Body: attribute values
- **Delete**: DELETE /v[apiVersion]/digital-objects/id

For example, to read an SFDO with ID of 13 using metadata version 0.1, the following request can be used: GET /v0.1/digital-objects/13.

While the backend is designed to be as independent as possible, as mentioned above the authentication is delegated to SLICES-RI global Authentication and Authorization Infrastructure (AAI) to ensure uniform user access across the entire infrastructure.

*4) Web portal:* The web portal allows user-friendly access to the MRS functionality. Users can search for objects (cf. Fig. 10), manage objects (if they have permissions), and access reporting facilities, such as dashboards. The web portal is accessible to all registered users.

The web portal is implemented as a single-page application (SPA) built using Angular 16. It interacts with the backend using the same REST JSON APIs as any other client would. Likewise, the web portal redirects the user to the global SLICES-RI AAI in order to obtain the access token (JWT) to add as a header to the backend API calls.

## VIII. Conclusions & Future Work

SLICES-RI is a unique initiative aimed at building a sustainable and scalable scientific instrument to support advanced experimentation in digital sciences. Its methodology is intent-based, anchored in the key scientific questions identified by diverse research communities. These questions are formalized through the concept of Blueprints — standardized references offered as a service to guide experimental design and execution.

We present SLICES-RI's current achievements, showcasing its capabilities through use cases in Post-5G and Federated Learning. We also invite the community to adopt our data management framework (MRS/DMI), a critical component designed to accelerate scientific progress and deepen understanding, while promoting reproducibility in computer systems research.

Looking ahead, several challenges remain that cannot be addressed without strong community engagement and collaboration. SLICES-RI is fully open — both in terms of access and source — and is intended to serve as a catalyst for cooperation. This commitment is already reflected in emerging partnerships aimed at fostering similar developments, including the DigitAfrica initiative and the OpenRIT 6G project in South America.

## References

[1] S. Fdida *et al.*, "SLICES, a scientific instrument for the networking community," *Computer Communications*, vol. 193, pp. 189–203, 2022.
[2] J. P. A. Ioannidis, "Why most published research findings are false," *PLoS medicine*, vol. 2, no. 8, p. e124, 2005.
[3] C. Collberg and T. A. Proebsting, "Repeatability in computer systems research," *Communications of the ACM*, vol. 59, no. 3, pp. 62–69, 2016.
[4] Y. Demchenko, S. Gallenmüller, S. Fdida, T. Rausch, P. Andreou, and D. Saucez, "Slices data management infrastructure for reproducible experimental research on digital technologies," in *2023 IEEE Globecom Workshops (GC Wkshps)*, pp. 1–6, IEEE, 2023.
[5] S. Gallenmüller, D. Scholz, H. Stubbe, and G. 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, p. 259–266, ACM, 2021.
[6] V. Q. Rodriguez, F. Guillemin, A. Ferrieux, and L. Thomas, "Cloud-RAN functional split for an efficient fronthaul network," in *2020 International Wireless Communications and Mobile Computing (IWCMC)*, pp. 245–250, 2020.
[7] D. Kefalas, N. Makris, S. Fdida, and T. Korakis, "Enhancing 5G V-RAN reliability via reactive CU-UP scaling," in *2025 European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit): Network Softwarisation (NET) (2025 EuCNC & 6G Summit - NET)*, (Poznan, Poland), p. 5.96, June 2025.
[8] F. Perez and B. E. Granger, "Project jupyter: Computational narratives as the engine of collaborative data science," *Retrieved September*, vol. 11, no. 207, p. 108, 2015.
[9] Y. Demchenko, S. Gallenmüller, S. Fdida, P. Andreou, C. Crettaz, and M. Kirkeng, "Experimental research reproducibility and experiment workflow management," in *15th International Conference on COMmunication Systems & NETworkS, COMSNETS 2023, Bangalore, India, January 3-8, 2023*, pp. 835–840, IEEE, 2023.
[10] I. Koren and R. Klamma, "The exploitation of openapi documentation for the generation of web frontends," in *Companion proceedings of the the web conference 2018*, pp. 781–787, 2018.
[11] N. Sakimura, J. Bradley, M. Jones, and E. Jay, "Openid connect dynamic client registration 1.0," *OpenID Specification. OpenID Foundation*, 2014.