Ref. Ares(2022)5084591 - 12/07/2022

slices
DS

HORIZON 2020
H2020 - INFRADEV-2019-3

D4.4 SLICES infrastructure articulation with NGI and international testbeds

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	February 28, 2022 (M18)
Submission Date	July 12, 2022 (M23)
Authors	Brecht Vermeulen (imec), Peter Van Daele (imec), Bartosz Belter (PSNC), Kostas Choumas (UTH), Yuri Demchenko (UvA), Cedric Crettaz (Mandat), Christian Perez (Inria), Lucas Nussbaum (Inria), Adrien Lebre (IMT).
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 deliverable is the first deliverable of task 4.3 (Relations with international testbeds) in WP4. It studies the relevant international testbeds and e-infrastructures related to advanced infrastructure and digital technologies, a.o. the Next Generation Internet (NGI) European testbeds and US infrastructures. It analyses interactions based on governance, technical cooperation, ad-hoc collaboration and guided collaboration (based on coordinated funding).
The following infrastructures were studied: Fed4FIRE – NGI, PlanetLab, EOSC, GENI, Emulab, CloudLab, Chameleon, FABRIC, PAWR, EMPOWER, PRACE.

It is advised that SLICES interacts with other infrastructures in the following ways:

- **Governance agreements**: it might be interesting to *define different levels of collaboration* at the governance level. In particular, SLICES is developing a collaboration effort at the international level, and in particular with the USA.
- **Technical cooperation**: all existing examples have technical cooperation while it seems that having *API standards (or pseudo standards in a specific large community) or using large community open-source software stacks seems the most beneficial.* It's important to note that larger infrastructures *do not use a single software stack for all sites* but they tend to agree on e.g., common APIs or common user tools to make it easy for the users. Cooperation or collaboration based on APIs (standards) and a reference architecture makes it possible to have future proof collaboration.
- **Ad-hoc collaboration**: SLICES should be open for collaboration but because of the size of SLICES, structured collaboration (e.g., through governance defined levels) seems the better choice. Ah-hoc collaboration might exist for very specific goals.
- **Guided collaboration**: by using funding from multiple sides, it is clearly shown that there is a higher interest in collaboration (e.g., EU-US), also this is done by the larger infrastructures.

Standardization can also be an enabler to interact with other infrastructures based on API definitions a.o. Based on the interaction/integration with existing infrastructures, this can also have an impact on the need to support one or more software stacks for the nodes.

Table of contents

List of figures

List of tables

1 Ways of infrastructure articulation

To start off, we should identify the different ways how infrastructures can relate to each other:

- **Governance agreements**: a common governance is possible or two or more governances can reach an agreement how to cooperate. This means that a strong link is possible, typically through contracts and maybe through funding/payment.
- **Technical cooperation or similarities**: if infrastructures use the same APIs or even the same software stack, an automatic relation is in place. If infrastructures are open, then users can easily use the different infrastructures without any learning curve. Multiple ways of technical cooperation can be identified as well:
    - users might even use a single account or tool to access multiple infrastructures.
    - or the users are even not aware that multiple infrastructures are in use.
    - The APIs are standardized and multiple infrastructures implement these APIs independently which makes a technical cooperation or tool re-use simple.
- **Ad-hoc collaboration**: infrastructures or users of infrastructures can always work together ad-hoc, e.g., by doing joined experiments (e.g., an experiment is run on two different infrastructures on a different moment in time and results are compared, or infrastructure is connected ad-hoc for an experiment, etc.).
- **Guided collaboration**: collaboration can also be stimulated, e.g., by funding organisations. E.g., project calls or travel funds are set up to stimulate the use of multiple infrastructures.

2 SLICES infrastructure and established methods

In this section we provide the reported SLICES features regarding the above dimensions of section 1. With respect to the governance agreements, the SLICES governance structure is based on the principle of a distributed RI with a Central Hub and a centralized governance. It is based on a classical type of governance structure of ESFRI RIs and other large research initiatives, catered to the specific needs of SLICES. The structure has three main levels with a Supervisory Board as the decision-making body of SLICES, a Coordination & Management Office (Executive Director and Thematic Directors) as the executive body in charge of the operational management of SLICES activities, and the Management Committee (with all Node Directors) that is in charge of the operational management coordination with all national nodes. The SLICES governance is complemented with an International Scientific Advisory Board, a User Committee, and several specific committees.

The overall governance of SLICES has been determined by closely monitoring existing actions around testing platforms for technologies that are included in the Digital Infrastructures field. As most of them are focusing from a specific point of view, they present different directions, depending on their application. For example, in a testbed federation, governance is managed at two levels and more loosely; a centralized governance and user management can enable users to use a single account for accessing different nodes of the infrastructure, while also local testbed access and testbed independence to implement their own mechanisms is allowed. As such, different infrastructures have different needs and thus different mechanisms for governance and mediating user access. The manner in which all of these have contributed to the overall SLICES governance are analyzed in the subsection further below.

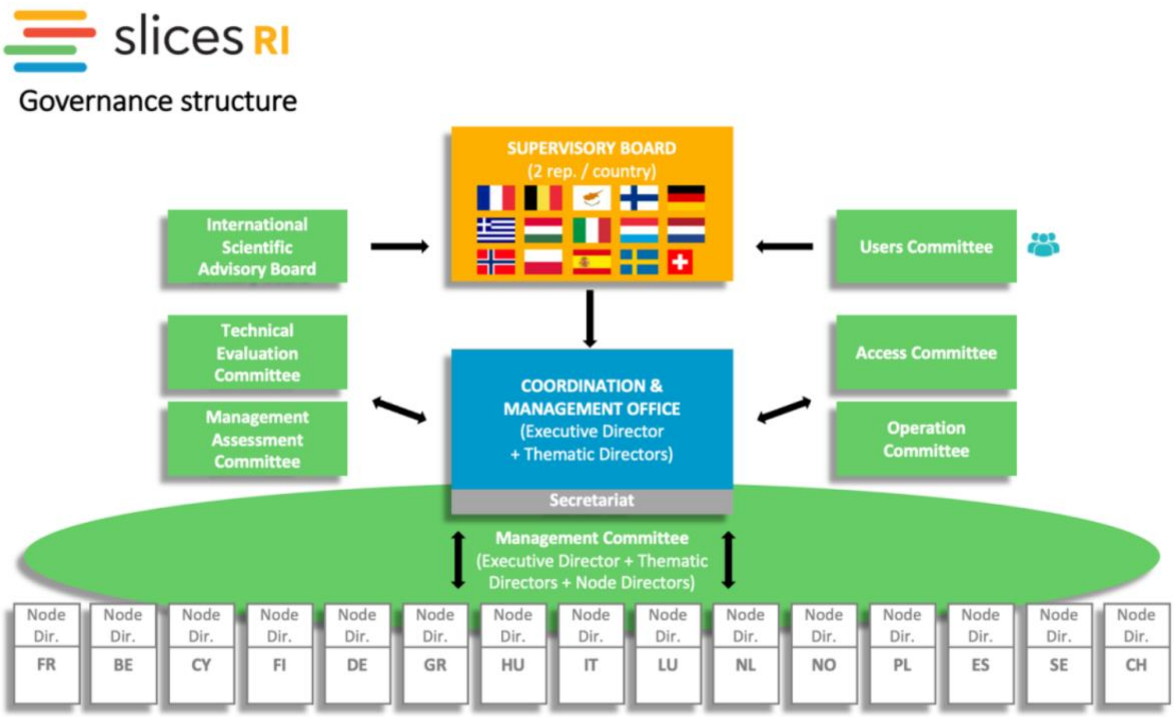

Figure 1: SLICES governance structure

3 Examples of infrastructure articulation within NGI, international testbeds and European e-infrastructures

In this section we list relevant big infrastructures or projects (not individual testbed nodes or sites).

3.1 Fed4FIRE – NGI

Fed4FIRE+ (https://www.fed4fire.eu) is a project under the European Union's Programme Horizon 2020, offering the largest federation worldwide of Next Generation Internet (NGI) testbeds, which provide open, accessible and reliable facilities supporting a wide variety of different research and innovation communities and initiatives in Europe, including the 5G PPP projects and initiatives.

It started in January 2017 and runs until June 2022. The Fed4FIRE+ project is the successor of the Fed4FIRE project (2012-2016). During 10 years, this federation of testbeds was built and funded by the EU. **The funding was specifically for building the federation, the testbeds were built and maintained by other funds (e.g., regional funding). The testbeds are also using different software stacks and APIs.** This is different from the other infrastructures mentioned in this document as these are built from the ground up around one software stack.

This is an example of infrastructure articulation that uses the following ways of infrastructure articulation:

- Governance agreement: through the grant agreements of the EU projects a governance agreement was defined. Nevertheless, centralized governance is not enforced across all the different entities, and thus each testbed can implement its own mechanisms within the Fed4FIRE larger framework. Some testbeds joined the federation without joining the project governance and this means there is less control on e.g., stability and maturity of such testbeds.

   - *Insights to SLICES governance:* the governance should be enforced enough on all testbeds, otherwise one loses maturity of the whole system. E.g., limit number of testbed technologies, prefer larger and more stable testbeds over smaller, testbeds should have a critical number of users, etc.
- Technical cooperation: Fed4FIRE agreed on using the same APIs for all infrastructures and as such users can use all infrastructures with a single account and a single tool. This was the main goal of the federation. This technical cooperation is now further brought to the newly set up ITU-T focus group on Testbed Federations (https://www.itu.int/en/ITU-T/focusgroups/tbfxg/Pages/default.aspx). A key component in the federation is also monitoring extensively the infrastructures (at API level and at full experiment cycle level): http://fedmon.fed4fire.eu . This monitoring info is also brought directly to the experimenter.

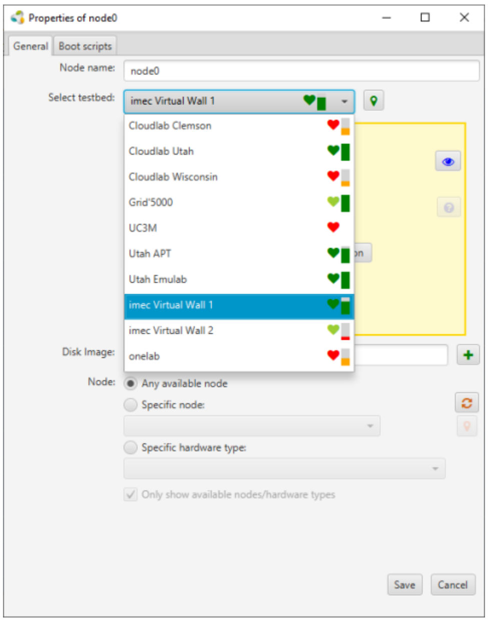

Figure 2: jFed GUI bringing the testbed monitoring (health, availability) information directly to the experimenter

- Ad-hoc collaboration: ad-hoc collaborations were set up during Fed4FIRE, e.g., with international testbeds such as GENI, Cloudlab, Chameleon, etc.
- Guided collaboration: also guided collaboration was set up between Fed4FIRE and other infrastructures, e.g., by organizing common summer schools with travel grants, open calls with funding in e.g., US and EU with partners in both continents.

3.2 PlanetLab

PlanetLab (https://en.wikipedia.org/wiki/PlanetLab) was an infrastructure consisting out of a group of computers available as a testbed for computer networking and distributed systems research. It was established in 2002 by Prof. Larry L. Peterson and Prof. David Culler (https://planetlab.cs.princeton.edu/history.html) and as of June 2010, it was composed of 1090 nodes at 507 sites worldwide. Each research project had a "slice", or virtual machine access to a subset of the nodes.

As such it was one of the predecessors of the GENI framework in the US.

- Governance agreement: there was 1 governance which parties could join. Accounts for using the testbed were limited to persons affiliated with corporations and universities that hosted PlanetLab nodes. As such, there was a legal binding between using the testbed and hosting

infrastructure. Since 2007 a second governance model with PlanetLab Europe was started. In 2020 PlanetLab central stopped while PlanetLab Europe still is up and running.
  - There was a governance agreement between PlanetLab central and Europe that users of one infrastructure could use the other infrastructure through federated accounts.
  - PlanetLab Europe also became one of the testbeds in the Fed4FIRE federation after adopting the necessary APIs.
  - *Insights to SLICES governance:* Planetlab was based on distributed governance, allowing users to login using their institutional accounts. Nevertheless, the process for single-sign on process to different geographically dispersed nodes has further evolved (and adopted by Fed4FIRE and other initiatives), and is a key feature in the overall SLICES governance model.
- Technically, the testbed was centrally driven and all nodes needed to be installed with the same pre-configured software (an image was created for each server. A local administrator downloaded and installed that image on their servers). The servers were further administered centrally by PlanetLab.
  - PlanetLab Europe also became one of the testbeds in the Fed4FIRE federation after adopting the necessary APIs.
  - The testbed software stack was used by some privately hosted/owned testbeds
- Ad-hoc collaboration: Measurement Lab (https://www.measurementlab.net/publications/mlab-founding-vision.pdf) is an example of a collaboration between OTI, PlanetLab, Google and academic researchers.

3.3 European Open Science Cloud (EOSC)

While the other infrastructure examples in this deliverable are targeted towards computer science research, EOSC is a lot broader and targets all scientific domains (https://eosc-portal.eu/). EOSC is the European federated scientific data infrastructure (https://eosc-portal.eu/) providing services to facilitate research data sharing between European Research Infrastructures and research institutions. The ambition of the European Open Science Cloud (EOSC) is to provide European researchers, innovators, companies and citizens with a federated and open multi-disciplinary environment where they can publish, find and reuse data, tools and services for research, innovation and educational purposes. For this purpose, EOSC established the market place portal: https://marketplace.eosc-portal.eu/ .

The EOSC enables a step change across scientific communities and research infrastructures towards:

- seamless data access and sharing.
- adoption of the FAIR (Findability, Accessibility, Interoperability and Reusability) data principles.
- publish and discover reliable reuse of infrastructure and research data services offered by European RIs, and all otherincluding digital objects, datasets, scientific workflows produced along the research life cycle (e.g., methods, software and publications).

The European Open Science Cloud (EOSC) ultimately aims to develop a 'Web of FAIR Data and services' for science in Europe upon which a wide range of value-added services can be built. These range from visualisation and analytics to long-term information preservation or the monitoring of the uptake of open science practices.

EOSC developed the EOSC Interoperability Framework (EOSC-IF) that defined 4 layers interoperability model to facilitate interoperability of infrastructures and services in the EOSC ecosystem, that is actually European Research Area, including layers: technical, semantic, organizational, and legal. Technical interoperability layer defines API for accessing and managing resources and services. Semantic interoperability layers define metadata profiles for services and data access and interaction. The EOSC portal has an API which allows to search, list and order the resources in the registry (https://providers.eosc-portal.eu/developers, https://marketplace.eosc-portal.eu/api_docs, see figure below). The API is required to support EOSC Provider and Resource metadata profiles. Cooperating infrastructures can use these APIs to link to EOSC.

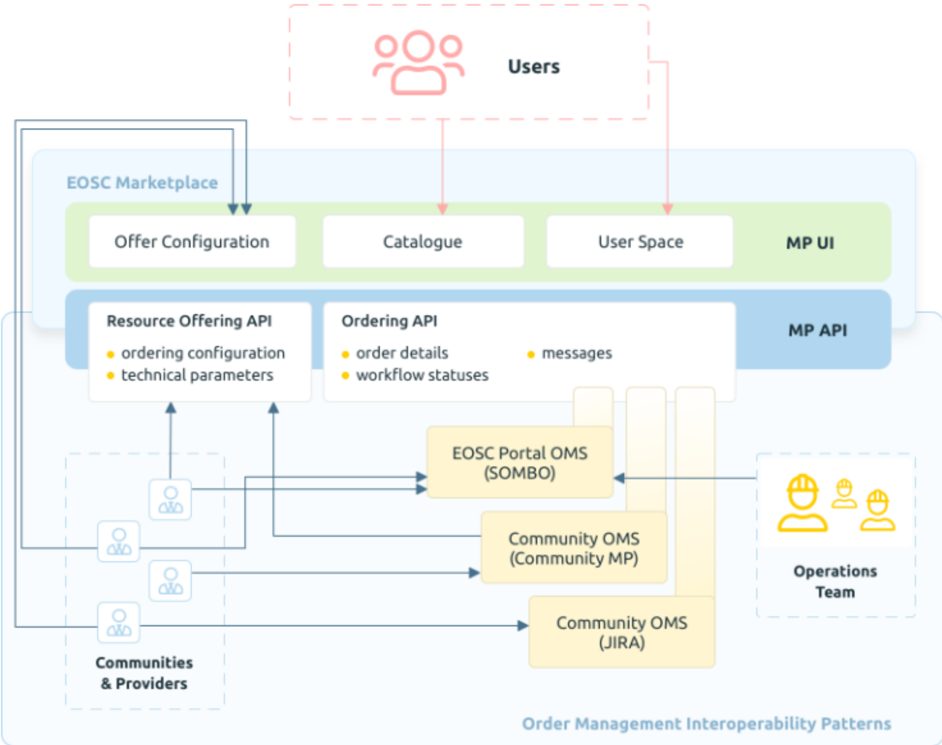

Figure 3: EOSC Marketplace APIs

- Governance: The EOSC tripartite collaboration is a concept of strategic coordination between the EU represented by the European Commission, the participating countries represented in the EOSC Steering Board, and the research community represented by the EOSC Association to resource and support the implementation of the EOSC ecosystem in Europe. This includes joint activities to enable and monitor the uptake of open science practices in Europe and to align relevant national and EU policies and investments to improve the production of FAIR research output that are "as open as possible, as closed as necessary".
  - *Insights to SLICES governance:* The EOSC model has directly impacted the manner in which data is managed and governed in SLICES. The respective guidelines for managing data produced over the infrastructure, as well as the complete metadata profiles and the creation of a SLICES entity who is in-charge of the data provisioning is directly fed into the SLICES governance model.
- Technical: SLICES Interoperability Framework adopts the EOSC Interoperability Framework and EOSC metadata profiles for Provider and Resource to ensure interoperability with EOSC and other RIs in the EOSC ecosystem.

- Collaboration: EOSC has a partnership with the European Commission (https://eosc.eu/partnership) and is open for everyone to "enable researchers to find, create, share and reuse all forms of digital knowledge – such as publications, data and software - leading to new insights and innovations, higher research productivity and improved reproducibility in science. To successfully support the digital transformation of science, this important Partnership must involve all actors – the research communities, our Universities and research institutions, service providers as well as our Member States."

3.4 Global Environment for Network Innovations (GENI)

GENI (https://geni.net) provides a virtual laboratory for networking and distributed systems research and education. It is well suited for exploring networks at scale, thereby promoting innovations in network science, security, services and applications. GENI was strongly influenced[1] by a number of precursors, notably PlanetLab, Emulab, DeterLab and ORBIT. The original concept of GENI was "a distributed Emulab/DeterLab with: PlanetLab's ability to host long-running services and experiments in virtualized environments; control of the inter-site networking; and ORBIT's ability to incorporate wireless nodes". GENI had two complementary lines of work: research on future architectures for the Internet and a robust infrastructure for experimenting with new and innovative infrastructural and application ideas. Funding started in about 2007-2008 and till today there is ongoing activity although some infrastructures have been phased out.

- Governance: The GENI project had the GENI Project Office (GPO) which steered and governed the GENI project, funded by the US NSF. As such a single governance organized the collaboration between the different funded tracks (e.g., Instageni, exogeni, orbit, …) and organized the GENI engineering conferences (GEC), tutorials, etc.
  - *Insights to SLICES governance:* The GENI centralized governance through the GPO is directly feeding into the SLICES governance model. Nevertheless, the processes for user management are altered, in order to address the different needs that the SLICES RI aspires to address.
- Technical collaboration: the GENI engineering conferences were open for everyone and (international) technical collaboration happened through the workshops at those meetings and the demo nights. Notably some of the APIs (Member Authority (MA) and Slice Authority (SA) APIs, https://geni-nsf.github.io/CommonFederationAPI/CommonFederationAPIv2.html) were developed in collaboration with Fed4FIRE. GENI was also very open for federation with other infrastructures implementing the same APIs (e.g., Fed4FIRE in Europe, but also Brazil, Japan, South-Korea).
- Ad-hoc collaboration: GENI was also open for ad-hoc collaboration, e.g., for international demos. The GENI Engineering Conferences played an important and successful role in this.
- Guided collaboration: NSF had bilateral agreements with several other countries and continents, e.g., matching open calls were organized with Europe, summer school travel grants were organized, etc.

3.5 Emulab

Emulab (https://www.emulab.net) is a testbed set up by the university of Utah in 2001. In many ways, the world's first Cloud and then and now the premier platform for controlled networking and

[1] Rick McGeer, Mark Berman, Chip Elliott, and Robert Ricci. 2016. The GENI Book (1st. ed.). Springer Publishing Company, Incorporated.

distributed systems experimentation. Emulab is a hardware-as-a-service cluster with elements devoted to controlled networking experimentation, primarily delay and traffic-shaping appliances.

- Governance: The University of Utah maintains a couple of related testbeds (Emulab, Cloudlab, Apt, Powder).
    - *Insights to SLICES governance:* Emulab has a centralized governance, similar to GENI, but to a smaller scale and with the big difference that Emulab is governed by a single organization (Univ of Utah). Similar insights on the models and the user management perspectives feed into the SLICES-RI model.
- Technical collaboration: the Emulab software stack is open (since about 2007) and can be used by other testbeds (e.g., DeterLab, instageni), and the University of Utah is also open for collaboration on this. E.g., the imec iLab.t testbeds (Virtual Wall, w-iLab.t, CityLab, https://doc.ilabt.imec.be) are running the Emulab software stack and imec has collaborated since 2006 on this software stack. The ProtoGENI APIs were one of the prototypes of the GENI APIs.
- Ad-hoc collaboration: University of Utah is open for ad-hoc collaborations, e.g., providing free accounts to their infrastructure.

3.6 CloudLab

CloudLab (https://cloudlab.us) is a flexible, scientific infrastructure for research on the future of cloud computing. Researchers use CloudLab to build their own clouds, experimenting with new architectures that will form the basis for the next generation of computing platforms.
CloudLab is built from the software technologies that make up Emulab and parts of GENI, so it provides a familiar, consistent interface for researchers.

- Governance: The University of Utah maintains a couple of related testbeds (Emulab, Cloudlab, Apt, Powder).
    - *Insights to SLICES governance:* CloudLab has a centralized governance, similar to GENI, but to a smaller scale and with the big difference that CloudLab is governed by a single organization (Univ of Utah). Similar insights on the models and the user management perspectives feed into the SLICES-RI model.
- Technical collaboration: the Emulab software stack is open (since about 2007) and can be used by other testbeds (e.g., DeterLab, instageni), and the University of Utah is also open for collaboration on this. CloudLab is further built on this software stack.
- Ad-hoc collaboration: University of Utah is open for ad-hoc collaborations, e.g., providing free accounts to their infrastructure.

3.7 Chameleon cloud

Chameleon (https://chameleoncloud.org) started in 2014 and is a large-scale, deeply reconfigurable experimental platform built to support Computer Sciences systems research. Community projects range from systems research developing new operating systems, virtualization methods, performance variability studies, and power management research to projects in software defined networking, artificial intelligence, and resource management.
To support experiments of this type, Chameleon supports a bare metal reconfiguration system giving users full control of the software stack including root privileges, kernel customization, and console access. While most testbed resources are configured in this way, a small amount is configured as a

virtualized KVM cloud to balance the need for finer-grained resource sharing sufficient for some projects with coarse-grained and stronger isolation properties of bare metal.

- Governance: Chameleon is an NSF funded project and the governance is formed by the project
    - *Insights to SLICES governance:* The Chameleon governance model is a centralized project based one. Insights on how equipment and users (APIs that the users can invoke for controlling the equipment) for infrastructure similar to Chameleon that exists in SLICES are mirrored in the SLICES-RI governance model.
- Technical collaboration: Unlike traditional Computer Science experimental systems which have overwhelmingly been configured by in-house infrastructures, Chameleon adapted OpenStack, a mainstream open-source cloud technology, to provide its capabilities. This has a range of practical benefits including familiar interfaces for users and operators, workforce development potential, leverage of contributions by a community 2,000 developers strong, and the potential to contribute to infrastructure used by millions of users (in particular, Chameleon team contributions to OpenStack include the Blazar component). Chameleon also adapted some of the GRID'5000 (EU testbed infrastructure) software stack components.

3.8 FABRIC

FABRIC (https://fabric-testbed.net, started in 2020) is a US-based nation-wide testbed supporting advanced networking, application and security research funded by the US National Science Foundation. It is a unique national research infrastructure to enable cutting-edge and exploratory research at-scale in networking, cybersecurity, distributed computing and storage systems, machine learning, and science applications. FABRIC is an everywhere programmable nationwide instrument comprised of novel extensible network elements equipped with large amounts of compute and storage, interconnected by high speed, dedicated optical links. It will connect a number of specialized testbeds (5G/IoT PAWR, NSF Clouds) and high-performance computing facilities to create a rich fabric for a wide variety of experimental activities.

- Governance: FABRIC is an NSF funded project and the governance is formed by the project. It will consist of 29 nodes/sites.
    - *Insights to SLICES governance:* The FABRIC governance model is a centralized project based one. Insights on how equipment and users (APIs that the users can invoke for controlling the equipment) for infrastructure similar to FABRIC that exist in SLICES are mirrored in the SLICES-RI governance model.
- Technical collaboration: The testbed code is open (https://github.com/FABRIC-testbed) but as far as we know the APIs and code are not yet used beyond the FABRIC project.
- Ad-hoc collaboration: as one of the key goals of FABRIC is to set up high-bandwidth interconnections between infrastructures, they are open for collaboration in setting up interconnectivity. They specifically have the FAB (FABRIC Across Borders) work package to set up interconnectivity and experiments across borders (with EU, Asia-Pacific and South-America).

3.9 PAWR

The Platforms for Advanced Wireless Research (https://advancedwireless.org/) program is enabling experimental exploration of new wireless devices, communication techniques, networks, systems, and

services that will revolutionize the nation’s wireless ecosystem while sustaining US leadership and economic competitiveness for decades to come.
Participating companies benefit by: helping to sustain US industry leadership; shaping design of research platforms; and securing cutting-edge research returns well in excess of initial investment. Participating communities benefit by: building core wireless capabilities through creative university partnerships; attracting government and corporate research funding and local wireless jobs; and utilizing advanced wireless capabilities to enhance city services and economic development.
PAWR is funded by NSF and a wireless Industry Consortium of 30 companies and associations (and this is different compared to the other infrastructures mentioned).

- Governance: The PAWR Project Office (https://advancedwireless.org/about-pawr-project-office/) is managing this $100 million public-private partnership to deploy and manage up to 4 city-scale research testbeds.
    - *Insights to SLICES governance:* The PAWR governance model is a centralized through the respective Project Office. Insights on user management (similar to the older GENI access) have fed in the governance model of SLICES.
- Technical collaboration: As far as we know the 4 platforms (POWDER in Salt Lake City, COSMOS in New York City, AERPAW in North-Carolina and ARA in Central Iowa) use different software stacks and are not integrated. The focus is more on the infrastructure and functionality than having a common user interface.
- Ad-hoc collaboration: PAWR is collaborating across the borders, e.g., with the EU EMPOWER project (see further).
- Guided collaboration: joint calls with e.g., EMPOWER.

3.10 EMPOWER

The EU CSA EMPOWER project (2018-2022) has the ambition to accelerate the joint development between the EU and the US of advanced wireless platforms targeting the new connectivity frontiers beyond 5G. EMPOWER targets the creation of a joint EU-US advanced wireless ecosystem for:

(i) bridging the relevant EU-US Wireless communities and stakeholders, such as scientific researchers, platform engineers, standardization experts, regulators, and product incubators.

(ii) developing a strategic EU-US collaboration agenda and supporting its execution ahead of worldwide competition for beyond 5G connectivity standards.

EMPOWER foresees twinning with the best researchers and practitioners involved in projects funded by USA, especially with entities participating in the PAWR programme (https://www.advancedwireless.org/). EMPOWER will provide instruments for inducing collaboration between ongoing and forthcoming 5G and beyond initiatives targeting at wireless networks experimentation on both ends of the Atlantic. Through the EMPOWER instruments we aim to create an efficient means for stimulating the mobility of ideas and people between European and similar American experimental wireless platform initiatives. We also aim at encouraging stronger collaboration between fundamental and experimental wireless researchers by making access to platform tools and data exchange simpler. EMPOWER instruments will also provide a wealth of information for global and regional standards and regulatory organizations (e.g., ITU-R, ETSI) and industry for a (e.g., NGMN). An important output of EMPOWER will also be in the form of recommendations on technologies and experimentation methodologies for future wireless experimentation objectives. This will assist in providing coordination between EU (FP9) and US NSF programmes for future individual and joint calls.

- Governance: the governance was formed through the CSA project
    - *Insights to SLICES governance:* The EMPOWER project has been a successful CSA, adopting several different boards in its governance affecting the collaboration of the project with external entities, and steering the project towards meaningful outcomes. The successful example of EMPOWER is continued in the SLICES-RI governance, through the participation of an International Scientific Advisory Board, a User Committee, and several specific committees, that monitor results and provide meaningful suggestions to the SLICES board.
- Technical collaboration: EMPOWER is a CSA project so is not focusing on technical developments/research but the EU advanced wireless platforms are running in the ICT-17 and ICT-19 calls, e.g.:
    - 5G-EVE: 5G European Validation platform for Extensive trails.
    - 5G-VINNI: An open large scale 5G end-to-end facility for KPI validation and verticals use case piloting.
    - 5GENESIS: An Open 5G Experimental Facility for Testing, KPI Validation and Showcasing.
- Guided collaboration: as described the CSA had as goal to collaborate with the US (PAWR) and set up e.g., joint calls.

3.11 PRACE

PRACE (https://prace-ri.eu/) is a research infrastructure for advanced computing in Europe. The mission of PRACE (Partnership for Advanced Computing in Europe) is to enable high-impact scientific discovery and engineering research and development across all disciplines to enhance European competitiveness for the benefit of society. PRACE seeks to realise this mission by offering world class computing and data management resources and services through a peer review process.
PRACE also seeks to strengthen the European users of HPC in industry through various initiatives. PRACE has a strong interest in improving energy efficiency of computing systems and reducing their environmental impact.
PRACE receives funding from three sources:

(i) The 7 leading-edge high-performance computing (HPC) systems (supercomputers) are funded and operated by five hosting countries (2016): France, Germany, Italy, Spain and Switzerland. During the initial period the four hosting partners (France, Germany, Italy and Spain) provided resources up to € 400 million over a five-year period (2010 – 2015). In 2016 a fifth Hosting Member, ETH Zurich/CSCS (Switzerland) opened its system via the PRACE Peer Review Process to researchers from academia and industry.

(ii) All 26 members of PRACE pay an annual fee and many of them provide national HPC resources as additional in-kind contributions.

(iii) The Implementation Phase of PRACE receives funding from the EU's Horizon 2020 Research and Innovation Programme (2014-2020) under grant agreement 730913.

- Governance: as a high-end research infrastructure, PRACE has set up a strong governance (https://prace-ri.eu/about/organisation/) as can be seen in the figure below.

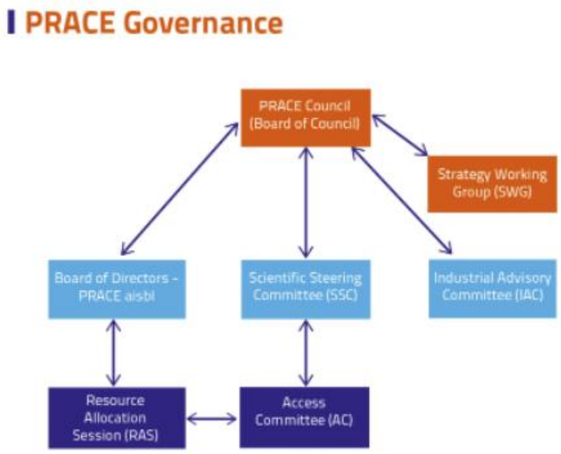

Figure 4: PRACE governance

Insights to SLICES governance: PRACE being one of the ESFRI projects, is highly influential to the manner that the SLICES governance is defined. Hence, SLICES has a similar board to which different committees and working groups report to. Moreover, the methodology for providing access to users for HPC resources is also closely monitored and influential to the manner that access to similar resources within SLICES-RI will be provided.

- Technical collaboration: as far as we could determine, PRACE is using a variety of software stacks or extensions, although it seems that most are based on the SLURM workload manager (https://en.wikipedia.org/wiki/Slurm_Workload_Manager) which is used on about 60% of the TOP500 supercomputers and as such is a widely contributed software stack.
- Ad-hoc collaborations: PRACE has a lot of collaborations and has also defined specific collaboration roles, such as memorandum of understanding, project partner, consortium member, representation in external board, and other collaborations). See https://prace-ri.eu/about/collaborations/ and https://prace-ri.eu/infrastructure-support/link-and-collaboration-with-other-e-infrastructure/
- Guided collaboration: PRACE is organizing collaborative calls with e.g., XSEDE, RIST and ICEI (https://prace-ri.eu/hpc-access/collaborative-calls/).

3.12 Summary

This table gives an overview of the ways of infrastructure articulation used by the mentioned infrastructures:

	Governance agreements	Technical cooperation	Ad-hoc collaboration	Guided collaboration
Fed4FIRE – NGI	X	X	X	X
PlanetLab	X	X	X	
EOSC	X	X	"must involve all actors – the research communities, our Universities and research institutions, service providers as well as our Member States"	
GENI	X	X	X	X
Emulab	X	X	X	
CloudLab	X	X	X	

Chameleon	X	X	X	
FABRIC	X	X	X	
PAWR	X	X	X	X
EMPOWER	X	Through ICT-17 and ICT-19 projects	X	X
PRACE	X	X	X	X

Table 1: Overview of ways of infrastructure articulation in existing infrastructures

The next table focusses on the centralized vs distributed governance.

	Centralized Governance	Distributed Governance
Fed4FIRE – NGI	++	+
PlanetLab	+++	
EOSC		+++
GENI	+++	
Emulab	+++	
CloudLab	+++	+
Chameleon	+++	
FABRIC	+++	+
PAWR		+++
EMPOWER	+++	
PRACE	+++	+

Table 2: Overview of ways of infrastructure articulation in existing infrastructures, regarding governance structure

4 Architecture and organization recommendations

4.1 Use cases and requirements for integrating SLICES-RI with external testbeds

This section provides several examples of basic scenarios for potential integration of SLICES-RI with external testbeds. Particular attention is paid to the requirements on technical and organizational aspects of such cooperation.

4.1.1 SLICES-RI experimenters require data available in the EOSC infrastructure

SLICES-RI will provide functionality to collect and maintain research data for the experiments carried out in the infrastructure. However, as explained in other project deliverables (D2.2 discusses this from the service point of view, D4.1 from the DMP point of view), this functionality will cover mostly short- and mid-term data preservation, offering SLICES-RI experimenters to manage their research data within the timeframe of their experiments. The general assumption taken by the SLICES-RI Consortium is that once the experiment is finished, the infrastructure will provide capabilities for the experimenters to analyze, extract and export research data to the EOSC infrastructure, where long-term preservation of data is provided.

In some cases, this functionality may not be sufficient for the experimenters, as they may require an access to archived research data for better experiment execution and testing reproducibility of an experiment they are carrying out. For this purpose, SLICES-RI should expose a dedicated interface towards EOSC, potentially using the existing EOSC developments, where the experimenters can search, retrieve and export research data from the EOSC infrastructure to SLICES-RI.

4.1.2 SLICES-RI experimenters require HPC resources available in the HPC Research Infrastructure

SLICES-RI is designed to provide a European-wide test-platform, equipped with advanced compute, storage and network components, interconnected by dedicated high-speed links. The computing part is the important element of the Research Infrastructure; however, the scope and character of this part is intentionally oriented on cloud/edge/fog computing and virtualization, rather than large scale HPC. Cloud/edge/fog computing is now an integral part of the future networking research and development. SLICES-RI will equip network researchers with capabilities to investigate different mechanisms in the clouds supporting innovation in the networking area (including 5G/6G and beyond). The rationale behind the decision of the SLICES-RI Consortium is also to not create another HPC Research Infrastructure, potentially competing with the existing infrastructures like PRACE or EGI but to offer advance HPC hardware to evaluate novel approaches or services before being transferred to production HPC centers.

The SLICES-RI experimenters may require in some cases an access to large scale computing power not available within the SLICES-RI. This scenario may cover the heavy-compute workloads as an outcome of experiments carried out in SLICES-RI to calculate advance AI models, extended simulation environments or even workloads requiring composition of multi-scale computations. This work should be carried out in third-party HPC-based facilities.

For this purpose, SLICES-RI should develop an interface towards PRACE, EGI or other HPC RIs, allowing the SLICES-RI experimenters to request, execute and retrieve the results of heavy-computing workloads. This work will require the technical analysis and agreements with the HPC RIs, both on organizational level (policies, governance) as well as on the technical level (APIs, authentication and authorization, etc.).

4.1.3 Third-party testbeds' experimenters require SLICES-RI resources

Within this deliverable several NGI testbeds have been identified and analyzed. It is important to ensure that the users of third-party NGI testbeds will be able to request and execute experiments in SLICES-RI, once they identify key software/hardware components available in SLICES-RI are critical to accomplish the planned experiment in the third-party testbed. The rules and policies for interconnecting SLICES-RI facilities and third-party testbeds should be developed and elaborated, as well as technical possibilities for linking all these resources should be investigated. As the work required to integrate with all identified NGI testbeds is enormous, it is recommended to proceed step by step as the specific needs of the research community emerge.

4.2 Technical recommendation on central vs distributed from a technical viewpoint

We have seen that in the existing infrastructure 4 big models can be identified:

The models mainly differ on whether sites can select the software stack of their choice to implement the API, and on who is in charge of maintaining the software stack. The models can be summarized as follows.

<table>
<tr><td>Model</td><td>Who selects the software stack ?</td><td>Who maintains the software stack ?</td></tr>
<tr><td>FED</td><td>Site (freely)</td><td rowspan="2">Site (possibly in collaboration with other sites sharing the same stack)</td></tr>
<tr><td>FED2SINGLE</td><td>Site (but encouraged to choose one of the official stacks)</td></tr>
<tr><td>FEDLIMITED</td><td>Site (but forced to choose one of the official stacks)</td><td>Software stack teams (1 team per software stack, funding TBD)</td></tr>
</table>

SINGLE	Central governance (only one software stack — possibly with various silos for some services to differenciate Cloud/HPC, IoT, wireless)	Software stack team (only one team, possibly distributed)

Apart from the governance model, the technical choices need also to be made and from this it seems that for new infrastructure a SINGLE model is very interesting, but when existing infrastructure need to be integrated or articulated e.g., a FEDLIMITED model might be more interesting and simpler to implement (this is allowing more than one software stack). This might be also relevant when multiple technologies need to be supported that are diverging quite much from each other. As SLICES aims to define future test platforms, following the ESFRI principle, a common reference architecture will be defined, allowing for the diversity that is necessary to achieve the ambition of the vision and mission.

5 Standardization

The standardization of existing APIs to integrate testbeds and to make them interoperable can provide some solutions on how to interact with other testbeds in the context of the SLICES Research Infrastructure. Indeed, some works on testbed federation and integration of different testbeds were recently undertaken in different SDOs such as ITU-T, ETSI and TM Forum. These advancements on the standardization of testbed APIs concern mainly the technical aspects and can be used not only in a testbed federation context, but also in other types of distributed testbed architecture.

The Recommendation ITU-T Q.4068 "Open application program interfaces (APIs) for interoperable testbed federations"[2] was published in August 2021 and presents the different APIs used by the Fed4FIRE+ testbed federation, namely the Aggregate Manager (AM) API, the Slice Authority (SA) API and the Member Authority (MA) API. These three APIs are commonly used by all the testbeds of the Fed4FIRE+ testbed federation, such IoT Lab, and by the GENI testbeds. By implementing such APIs, the SLICES Research Infrastructure can directly access the resources provided by the Fed4FIRE+ testbeds and GENI testbeds.

Furthermore, the Recommendation ITU-T Q.4068 describes a generic reference model for testbeds federations based on the ETSI TC INT GANA model[3]. The GANA (Generic Autonomic Networking Architecture) model permits to create an ecosystem around the testbeds enabling automation activities like resource discovery among the different stakeholders (testbed suppliers, researchers, etc.). This model allows the specialization of testbed domain. For instance, different kinds of testbeds can be integrated in one entity or organization; IoT testbeds, wireless network testbeds, RAN testbeds, MEC testbeds can be incorporated and consequently made interoperable inside the same organization.

[2] ITU-T Q.4068, *Open application program interfaces (APIs) for interoperable testbed federation,* https://handle.itu.int/11.1002/1000/14765 [Last accessed 12 July 2022]

[3] Autonomic network engineering for the self-managing Future Internet (AFI); Generic Autonomic Network Architecture; Part 2: An Architectural Reference Model for Autonomic Networking, Cognitive Networking and Self-Management, ETSI TS 103 195-2 V1.1.1, https://www.etsi.org/deliver/etsi_ts/103100_103199/10319502/01.01.01_60/ts_10319502v010101p.pdf [Last accessed 12 July 2022]

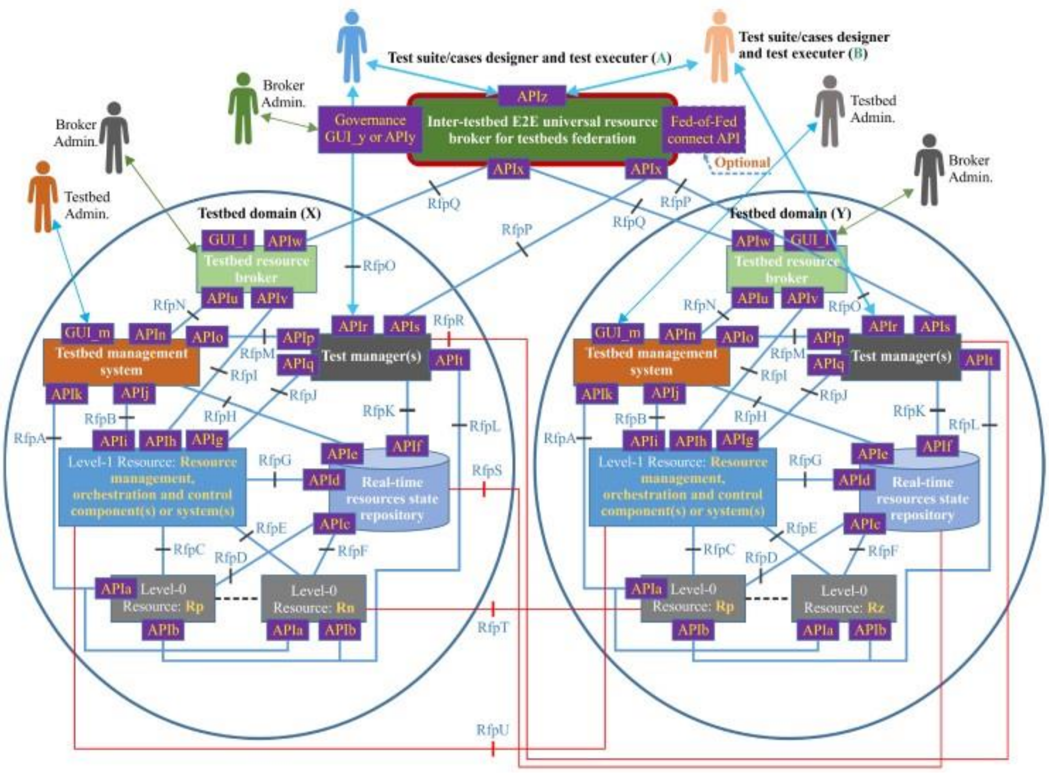

Figure 5: Generic federated testbed model (source: Recommendation ITU-T Q.4068, figure 1)

New use cases and business models, notably for the Testbed as a Service (TaaS), are currently developed around the Recommendation ITU-T Q.4068 and the included generic reference model in the dedicated ITU-T Focus Group on Testbeds Federations for IMT-2020 and beyond (FG-TBFxG)[4]. The work currently done in this Focus Group is based on the Recommendation ITU-T Q.4068 and concerns the development of APIs related to the testbeds federations reference model, including their specifications and standardization. The Focus Group also serves as a platform to exchange views on concepts related to testbeds to make them interoperable through dedicated APIs. The Focus Group is open to all SDOs, fora and organizations. The final objective of the Focus Group is to delivered a set of standardized APIs allowing the interconnection and the interoperability of testbeds across the world. These APIs can be potentially used by the SLICES Research Infrastructure to interact with other international research infrastructures or testbeds.

Other developments about the testbeds were realized by TM Forum. Indeed, this global industry association group the communications service providers (CSPs) has recently designed Open APIs to execute tests in a dedicated development and test environment. Basically, three TM Forum Open APIs are available from the TM Forum GitHub[5]:

- TMF705 Test Environment Management API: This API allows the management of a whole test environment, corresponding to a testbed.
- TMF708 Test Execution API: The API is focused on the entire life cycle of an experiment, including of course its execution.
- TMF913 Test Component API Suite: This suite of APIs encompasses in fact the TMF705 Test Environment Management API and the TMF708 Test Execution API. The set of APIs permits to support DevOps not only in the software development, but also the related IT operations.

[4] ITU-T Focus Group on Testbeds Federations for IMT-2020 and beyond (FG-TBFxG), https://www.itu.int/en/ITU-T/focusgroups/tbfxg/Pages/default.aspx [Last accessed 12 July 2022]
[5] TM Forum GitHub: https://github.com/tmforum-rand (requires a TM Forum account to access all the projects) [Last accessed 12 July 2022]

These TM Forum Open APIs are based on models defined by TM Forum and commonly used by the communications service providers and the telecommunications operators. This selection of TM Forum Open APIs related to test environments could be potentially reused to interact with testbeds deployed by the large TM Forum community.

6 Conclusion: SLICES infrastructure articulation with other testbeds

From these insights in collaborations and agreements, it seems straightforward that also SLICES has to interact with other testbeds and e-infrastructures, in the EU and beyond:

- **Governance agreements**: it might be interesting to *define different levels of collaboration* at the governance level. E.g., if infrastructures would like to join or connect to SLICES, they need to adhere to the same APIs, maturity and functionality levels, etc. (which need to be defined).
    - As we have also defined in D3.4, SLICES has adopted a centralized governance structure, which is based on a classical type of governance structure of ESFRI research infrastructures and other large research initiatives, catered to the specific needs of SLICES. Most of the aforementioned testbeds also follow some form of centralized approach for their governance, such as PlanetLab, GENI, Emulab, Cloudlab, Chameleon, FABRIC, EMPOWER (examples where the testbed infrastructure is built after defining the governance). Last but not least, the most important lesson learnt is from the Fed4FIRE successful example, which addressed the federation between various existing testbeds (where the testbeds themselves had already a local governance), but its further evolution needs a more centralized approach that will be the cornerstone of the SLICES governance approach. An important decision/question is also how to interact with existing infrastructure (with an existing governance).
- **Technical cooperation**: all examples have technical cooperation while it seems that having *API standards (or pseudo standards in a specific large community) or using large community open-source software stacks seems the most beneficial.* It's important to note that larger infrastructures (Fed4FIRE, GENI, PAWR, PRACE) *do not use a single software stack for all sites* but they tend to agree on e.g., common APIs or common user tools to make it easy for the users. Cooperation or collaboration based on APIs (standards) makes it possible to have future proof collaboration. E.g., Emulab created the ProtoGENI prototype API, which was adopted by GENI and further extended and improved to build the instageni and exogeni infrastructures. These APIs were also adopted by Fed4FIRE and further extended with extra APIs by GENI and Fed4FIRE. Future generation of platforms, including in the US will seek for adopting a reference architecture allowing for the diversity necessary to support the vision of such platforms.
- **Ad-hoc collaboration**: all projects/infrastructures are open for collaborations; some are more organized than others. SLICES should be open for collaboration but because of the size of SLICES, structured collaboration (e.g., through governance defined levels) seems the better choice.
- **Guided collaboration**: by using funding from multiple sides, it is clearly shown that there is a higher interest in collaboration (e.g., EU-US), also this is done by the larger infrastructures.

As discussed, standardization can also be an enabler to interact with other infrastructures based on API definitions a.o.

Based on the interaction/integration with existing infrastructures, this can also have an impact on the need to support one or more software stacks for the nodes.

slicesDS