



# Fenix: Distributed e-Infrastructure Services for EBRAINS

Sadaf Alam<sup>1</sup>, Javier Bartolome<sup>2</sup>, Sanzio Bassini<sup>3</sup>, Michele Carpenè<sup>3</sup>,  
Mirko Cestari<sup>3</sup>, Frederic Combeau<sup>4</sup>, Sergi Girona<sup>2</sup>, Stefano Gorini<sup>1</sup>,  
Giuseppe Fiameni<sup>3</sup>, Björn Hagemeier<sup>5</sup>, Andreas Herten<sup>5</sup>, Nikoleta Kiapidou<sup>2</sup>,  
Wouter Klijn<sup>5</sup>, Dorian Krause<sup>5</sup>, Jacques-Charles Lafoucrière<sup>4</sup>, Cerlane Leong<sup>1</sup>,  
Thomas Leibovici<sup>4</sup>, Thomas Lippert<sup>5</sup>, Colin McMurtrie<sup>1</sup>, Pavel Mezentsev<sup>5</sup>,  
Anne Nahm<sup>5</sup>, Boris Orth<sup>5</sup>, Dirk Pleiter<sup>5(✉)</sup>, Thomas Schulthess<sup>1</sup>,  
Benedikt von St. Vieth<sup>5</sup>, Debora Testi<sup>3</sup>, and Gilles Wiber<sup>4</sup>

<sup>1</sup> CSCS, Swiss National Supercomputing Centre, 6900 Lugano, Switzerland

<sup>2</sup> Barcelona Supercomputing Center, 08034 Barcelona, Spain

<sup>3</sup> CEA DAM, 91680 Bruyères-le-Châtel, France

<sup>4</sup> CINECA, 40033 Casalecchio di Reno, Italy

<sup>5</sup> JSC, Forschungszentrum Jülich, 52425 Jülich, Germany

pleiter@kth.se

**Abstract.** The Human Brain Project (HBP) (<https://humanbrainproject.eu/>) is a large-scale flagship project funded by the European Commission with the goal of establishing a research infrastructure for brain science. This research infrastructure is currently being realised and will be called EBRAINS (<https://ebrains.eu/>). The wide ranging EBRAINS services for the brain research communities require diverse access, processing and storage capabilities. As a result, it will strongly rely on e-infrastructure services. The HBP led to the creation of Fenix (<https://fenix-ri.eu/>), a collaboration of five European supercomputing centres, who are providing a set of federated e-infrastructure services to EBRAINS. The Fenix architecture has been designed to uniquely address the need for a wide spectrum of services, from high performance computing (HPC) to on-demand cloud technologies to identity and access federation, for facilitating ease of access and usage of distributed e-infrastructure resources. In this article we describe the underlying concepts for an audience of computational science end-users and developers of domain-specific applications, workflows and platforms services. To exemplify the use of Fenix, we will discuss selected use cases demonstrating how brain researchers can use the offered infrastructure services and describe how access to these resources can be obtained.

**Keywords:** Fenix · Human brain project · Distributed e-infrastructures

## 1 Introduction

Today, advances in addressing grand challenges depend on the ability of researchers coming from different geographic regions to effectively collaborate

and having flexible access to distributed e-infrastructure services. Fenix supports the former and provides the latter as it enables researchers, e.g., to collaborate on data curation, aggregation and sharing by providing federated storage sources as well as to use high-capability resources like High Performance Computing (HPC) systems.

The brain research community is a diverse community that applies a large variety of methods. It is thus not surprising that the requirements concerning e-infrastructure services are rather diverse. While some teams need massively-parallel HPC systems for large-scale simulations, others are producing extreme-scale data sets and employ advanced and potentially compute-intensive data analysis techniques. Given the size of the data sets, data analysis and making these data sets available to the wider community needs to be done in such a way that data transport can be avoided. This does not only help reducing costs or improving performance, in a number of cases this is mandatory given the involved amount of data. Yet other research teams perform computational and data analysis tasks requiring compute resources at smaller scale, but need to do this in an interactive manner.

Fenix is creating an infrastructure layer comprising services that are federated in a rather lightweight fashion. It is designed such that quite different types of compute services ranging from HPC to Cloud as well as different types of data repositories can be integrated. This initiative is realised by five European supercomputing centres, namely BSC in Spain, CEA in France, CINECA in Italy, CSCS in Switzerland, and JSC in Germany. Fenix is organised in such a way that other resource providers may join in the future.

This paper is organised as follows: In the next section we introduce the general concepts that led to the architecture of Fenix. In Sect. 3 we provide details on the Fenix services and discuss in Sect. 4 how EBRAINS services can make use of these. In Sect. 5 we describe how resources are allocated to users from HBP, before providing a summary and outlook in Sect. 6.

## 2 Fenix Concept

The current architecture of Fenix is based on the general consideration that a clear separation between an infrastructure service layer and a platform service layer is beneficial. Such a layered approach is commonly used for creating Cloud infrastructures (see, e.g., [2]), where the terms Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) are widely used.

For Fenix we prefer to use the terms infrastructure service layer and platform service layer. The *platform service layer* encompasses all services that are specific for a given research domain. They are not necessarily useful for other domains or would require significant adaptations. A typical example are web-based portals, such as the EBRAINS Collaboratory<sup>1</sup>. While such portals are needed for almost any research infrastructure, their organisation is highly domain specific. The

---

<sup>1</sup> <https://wiki.ebrains.eu>.

*infrastructure service layer* includes a set of services that allow implementing these platform services and are sufficiently generic for being useful for different research communities. One example are machines for deploying any of the aforementioned portal services, which are typically offered in a virtualised environment. Using Virtual Machine (VM) technology allows for better exploitation of the hardware resources as a larger number of VMs, which typically only need the resources of a few CPU cores, can be deployed within a single physical machine.

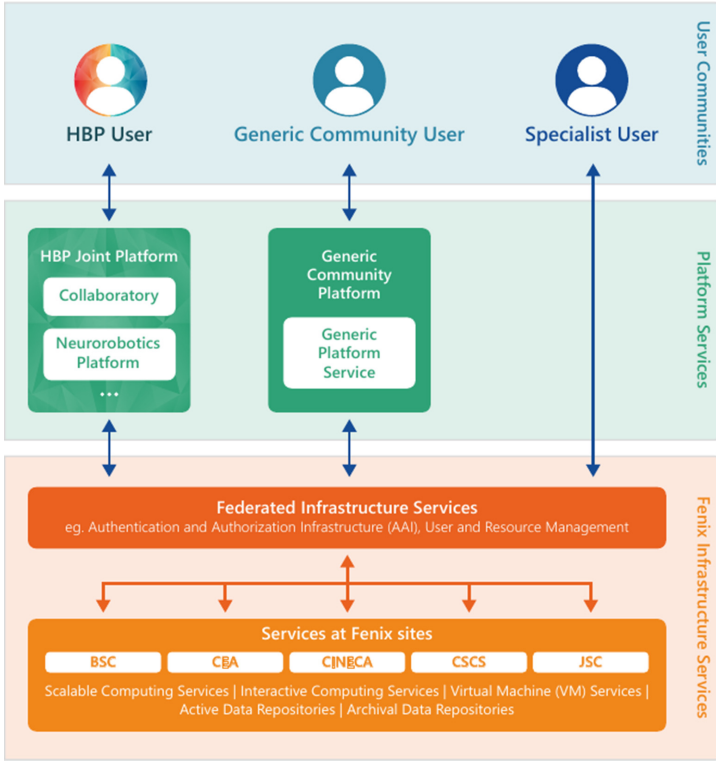
The infrastructure services are organised such that they can be provided by multiple, geographically distributed resource providers. While this approach adds the complexity related to the federation of these services, the approach has a number of important benefits, which we will discuss below.

Most of the end-users are not expected to use the Fenix infrastructure services directly, but rather connect to the platform services that are deployed on top of these infrastructure services as shown in Fig. 1. For specialist users there is, however, the option to directly access the infrastructure services. An example are users performing simulations on massively-parallel HPC systems, who typically directly access these systems to compile and execute their simulation applications.

The layered approach has multiple benefits. In general, a layered approach and the resulting separation of concerns helps to manage complexity. From the perspective of the platform service providers, the abstraction of an infrastructure service layer can help improving sustainability and performance due to the distributed nature of the infrastructure service layer, involving multiple infrastructure resource providers. Resource providers can be replaced, for instance when funding conditions change, or the number of resource providers could be changed according to the needs of the platform service layer. Furthermore, platform service providers are enabled to improve on resilience by replicating their services over multiple sites. Another benefit of a distributed infrastructure is improved data locality. With a larger number of infrastructure resource providers, the probability that storage resources are available in geographic proximity of the data source increases. From the perspective of the infrastructure service providers, the layered approach has the benefit that it allows the consolidation of their service offerings when supporting multiple science communities. Finally, it creates opportunities for improving the utilisation of the offered hardware resources.

### 3 Fenix Compute and Data Services

In this section we provide an overview of the current service portfolio offered by Fenix, which was developed on the basis of an analysis of today's needs of the brain research communities. With other communities starting to use Fenix, the current portfolio of services is anticipated to change.



**Fig. 1.** Overview of the Fenix architecture as described in Sect. 2. Details on the infrastructure services are provided in Sect. 3.

The *Scalable Compute Services (SCC)* abstract large-scale computing resources. These are HPC systems with a larger number of compute nodes with 1–2 CPUs and possibly additional compute accelerators like GPUs. SCC services can be used for running highly parallel simulation applications, but are also suitable for data analysis tasks, involving extreme-scale data sets. SCC resources are managed by a batch queuing system, which schedules jobs such that hardware utilisation is optimised.

As an increasing demand for interactive access to compute resources is observed, Fenix introduces *Interactive Compute Services (IAC)* services. These allow end-users to obtain ad-hoc access to single compute nodes where interactive frameworks like Jupyter are offered. Typical usage scenarios are interactive analyses, visualisations, and steering of simulations running on SCC.

*VM* services offer access to on-demand virtualised machines. The prime use case for this service is the deployment of platform services running in a “24/7” mode, for instance web-based portal services.

To cope and comply with different and in parts incompatible needs and requirements, Fenix introduces two classes of data repositories. An *Archival Data*

*Repository (ADR)* is a storage system for long-term storage of data objects in a shareable manner. Such data repositories must therefore feature a standardised interface with easy to install clients and allow for federation. Mechanisms supporting flexible and fine-grained access control are another important feature. Fenix decided for the widely used Cloud object storage interface Swift<sup>2</sup>.

Unlike an ARD, an *Active Data Repository (ADR)* is not a federated data repository that is relatively openly accessible from outside a data centre. An ACD is meant to be used for storing (copies of) private data sets and will typically offer 10 – 100× more bandwidth as well as significantly lower latency. Such features are important for data repositories connected to SCC services. A typical implementation of an ACD is based on a parallel file system with a POSIX interface like Lustre<sup>3</sup> or Spectrum Scale<sup>4</sup>.

Both types of data repositories will be connected through a *Data Mover* service that will allow to asynchronously copy or move data back and forth.

It is important to note that the different services have different security requirements. Some of them, like SCC, IAC and ACD, are realised in an *HPC environment* with tightly restricted access policies. VM services are deployed in a *Cloud environment* with an open connectivity. This will, e.g., allow users unknown to the data centre with weak or no credentials to connect to the platform services deployed on such resources. An ARD is connected to both environments and thus can serve as a bridge between both worlds. Other connections between services deployed in different environments are subject to negotiations to identify the right balance between the realisation of advanced workflows and security concerns.

Except for these restrictions, Fenix allows to combine the use of different services at different locations within a single project. This is a significant advantage compared to similar service offerings in Europe. Achieving this depends on the following prerequisites:

- All services must be integrated into a single Authentication and Authorisation Infrastructure (AAI) such that a user can connect with the same credentials to any service offered at any Fenix site.
- Resource management must be centralised such that both, a Fenix user as well as a Fenix resource provider, have an overview of the resources that are still available or have already been consumed.
- For coherent management of access control, a central service is needed that makes the necessary attributes available.

At the time of writing this article, a first version of the AAI is being put in operation, while the central resource management and attribute service, which is called FURMS, is still under development.

While the Fenix service portfolio is provided using general-purpose hardware technologies, the concept described in this section also allows for provisioning

---

<sup>2</sup> <https://wiki.openstack.org/wiki/Swift>.

<sup>3</sup> <https://lustre.org/>.

<sup>4</sup> <https://www.ibm.com/products/spectrum-scale>.


of infrastructure services, using special-purpose hardware solutions. Within the HBP there are, e.g., ongoing efforts to integrate neuromorphic computing services that are provided on the BrainScales [8] and SpiNNaker [4] systems at the Universities of Heidelberg and Manchester, respectively.

## 4 Selected EBRAINS Services

In this section we introduce a selected set of EBRAINS services and discuss how these can make use of Fenix services.

The EBRAINS *Brain Simulation Platform* comprises a suite of software tools and workflows for collaborative brain research that allow researchers to reconstruct and simulate detailed models of brain areas. This includes, e.g., simulators like NEST [5] and Neuron [7] as well as the neurophysiology data analysis tools package Elephant [3]. A simple workflow using Fenix services is shown in Fig. 2:

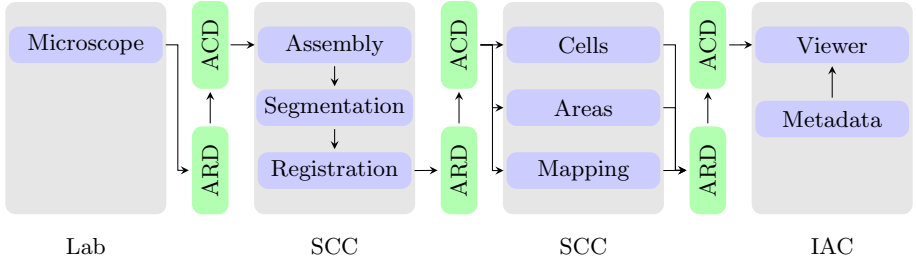
1. The input model data is assumed to be stored in ARD #1 and is copied to an ACD from where it is accessible to the SCC service.
2. The simulations are executed using the SCC service, which reads the input data from and writes the output data to the ACD.
3. After completing the simulation, the final data products can be published by copying the data to ARD #2.



**Fig. 2.** Example for a brain simulation flow using Fenix services.

Next, we consider a more complex example related to the EBRAINS *Brain Atlases*. EBRAINS aims to provide access to a new generation of 3-dimensional reference atlases of the human and rodent brain, which are defined at different scales and modalities. These atlases are based on histological data obtained from brain images (see, e.g., [1]). A complex workflow is required to first analyse and interpret the images as well as to integrate this data (see, e.g., [6]) and to later make the Brain Atlas as well as primary and secondary data products available to others. A possible realisation of such a workflow is shown in Fig. 3, which can be mapped to Fenix services as follows:

1. The primary data products are generated in a lab and stored in an ARD.
2. SCC services allow processing of extreme-scale data sets as they occur in the case of images with a resolution of  $\mathcal{O}(1\mu\text{m})$ , and facilitate the use of compute-intensive data analysis steps. To allow for fast access to the data, it will typically be staged from an ARD to an ACD. The resulting data products can be published after writing them into an ARD.
3. Multiple analysis steps using SCC services may follow.
4. Final data products may be explored interactively using IAC services.



**Fig. 3.** Schematic view of a possible workflow for creating a Brain Atlas using Fenix services.

## 5 Resource Allocation

Part of the resources provided by Fenix are dedicated to HBP research and EBRAINS services as well as related research projects. A so-called programmatic access model allows to provide these resources to a research community (here the brain research community represented by HBP), with the latter being responsible for the allocation of the resources to projects proposed by researchers from that community.

The HBP allocates the resources based on a peer-review mechanism that follows principles established by PRACE<sup>5</sup>. Such a mechanism is widely used for making HPC resources available as it helps to ensure expensive resources being used for excellent science. The principles mandate, among other requirements, the peer-review process to be transparent and clear to all relevant stakeholders. Furthermore, the process must be fair such that all proposals are evaluated solely on merit and potential high impact on European and international science and economy.

Applicants interested in using Fenix resources for brain research can, at any time, submit a proposal. After a technical review by Fenix resource providers,

<sup>5</sup> <https://prace-ri.eu/hpc-access/project-access/project-access-the-peer-review-process/>.

the EBRAINS Infrastructure Allocation Committee (IAC) is responsible for conducting or managing the scientific assessment in case of small- or large-scale resource requests, respectively. Based on the outcome of the review, the IAC can in the case of small-scale projects decide itself on whether to approve or reject a proposal. In the case of large-scale projects, the IAC prepares a decision-making proposal for the Directorate of the HBP.

## 6 Summary and Outlook

Fenix is an initiative that is realising a broad set of federated infrastructure services. The approach is based on a generic concept that aims for a separation of infrastructure and platform services. While the former are generic and of use for a variety of research communities, the latter are research domain specific. The approach allows research communities to establish distributed research infrastructures adapted to their needs. The brain research community is the key driver for Fenix and two examples of how this community can leverage services and resources from Fenix have been discussed. Similar efforts towards IT-based, distributed research infrastructures can, however, also be observed for other science communities.

The HBP has established mechanisms for allocating resources offered by Fenix, which is open for researchers from HBP but also for brain researchers at large. Other scientists can also apply for Fenix resources through regular calls for proposals managed by PRACE.

**Acknowledgements.** Funding for the work is received from the European Commission H2020 program under Specific Grant Agreement No. 800858 (ICEI). We would like to thank Timo Dickscheid (Forschungszentrum Jülich (FZJ), INM-1) for many discussions providing insight into the workflow for creating a brain atlas from brain imaging data.

## References

1. Amunts, K., Zilles, K.: Architectonic mapping of the human brain beyond brodmann. *Neuron* **88** (2015) <https://doi.org/10.1016/j.neuron.2015.12.001>
2. Badger, L., Grance, T., Patt-Corner, R., Voas, J.: Cloud Computing Synopsis and Recommendations. Tech. rep, NIST (2012)
3. Denker, M., Yegenoglu, A., Grün, S.: Collaborative HPC-enabled workflows on the HBP laboratory using the elephant framework. In: *Neuroinformatics 2018*, p. P19 (2018). <https://doi.org/10.12751/incf.ni2018.0019>
4. Furber, S.B., Galluppi, F., Temple, S., Plana, L.A.: The SpiNNaker project. *Proc. IEEE* **102**(5), 652–665 (2014). <https://doi.org/10.1109/JPROC.2014.2304638>
5. Gewaltig, M.O., Diesmann, M.: NEST (NEural Simulation Tool). *Scholarpedia* **2**(4), 1430 (2007)
6. Lebenberg, J., et al.: A framework based on sulcal constraints to align preterm, infant and adult human brain images acquired in vivo and post mortem. *Brain Struct. Funct.* **223**(9), 4153–4168, December 2018



7. Lytton, W.W., Seidenstein, A.H., Dura-Bernal, S., McDougal, R.A., Schürmann, F., Hines, M.L.: Simulation neurotechnologies for advancing brain research: parallelizing large networks in NEURON. *Neural Comput.* **28**(10), 2063–2090 (2016). [https://doi.org/10.1162/NECO\\_a\\_00876](https://doi.org/10.1162/NECO_a_00876)
8. Schmitt, S., et al.: Neuromorphic hardware in the loop: training a deep spiking network on the BrainScaleS wafer-scale system. In: 2017 International Joint Conference on Neural Networks (IJCNN), pp. 2227–2234 (2017). <https://doi.org/10.1109/IJCNN.2017.7966125>

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

