

Fenix: A Pan-European Federation of Supercomputing and Cloud e-Infrastructure Services

BY SADAF R. ALAM, JAVIER BARTOLOME, MICHELE CARPENE, KALLE HAPPONEN, JACQUES-CHARLES LAFOUCRIERE, AND DIRK PLEITER

TO ADDRESS INCREASING demands of various research communities for computing and storage services, six leading European supercomputing centers began harmonizing and federating their e-infrastructure services portfolio with the goal of supporting a variety of science and engineering communities. Barcelona Supercomputing Centre (BSC) in Spain, France's Commissariat à l'énergie atomique et aux énergies alternatives (CEA), Italy's Supercomputing Centre (CINECA), Finland's Supercomputing Centre (CSC), the Swiss National Supercomputing Centre (CSCS), and Juelich Supercomputing Centre (JSC) in Germany, have aligned high-end computing and storage services to facilitate the creation of the Fenix Research Infrastructure, which has been making resources available at scale to research communities since 2018.²

Characterized by different types of data repositories, scalable supercomputing systems, and private cloud instances, the Fenix portfolio is complemented by a federated identity and access management system.³ Presently, a diverse

portfolio of services are available for HPC, AI, and ML, and cloud computing applications, free of charge (<https://fenix-ri.eu/access>). Evaluation of the applications follows the peer review principles established by PRACE (<https://prace-ri.eu/>). The Fenix objective is to serve science and engineering domains that strongly benefit from diverse e-infrastructure services for their collaborative research and data sharing. It therefore leverages national, European, and international funding programs to realize the compute, storage, and network resources sustaining the e-infrastructure services. There are similar national programs such as U.S. NSF XSEDE (<https://www.xsede.org/>). However, Fenix introduces unique aspects: First, it defines a federated research e-infrastructure architecture for leadership-class supercomputing resources providers, transcending national boundaries; and second, it offers a uniform, federated identity, and access management solution.

A Co-Designed Solution for Science and Scientists

Development of Fenix services and the underlying technical solutions has been an iterative, co-design




Figure 1. The Fenix consortium, and its federated and e-Infrastructure services portfolio.

process, which was initially driven by the Human Brain Project (HBP)—a flagship venture currently funded by the European Commission for a period of 10 years. Fenix has been facilitating the design, implementation, and operations of domain-specific (neuroscience) platform services. The need to federate services arose from a collaboration between scientists working on, for example, a Brain Atlas, an HBP platform service that requires integration of data from a variety of research teams

throughout Europe, and beyond.⁴ Various neuroscience workflows have in common the need of being able to collect data at the edge (that is, instruments running measurements), moving it to a nearby or affiliated datacenter, and making it available for further compute-intensive processing and integration with other datasets. The latter may come from geographically distributed datacenters.


Figure 1 highlights the concept of Fenix that has been co-designed with neu-

rosience and other similar use cases, for instance, the Materials Cloud, which is a platform designed to enable open and seamless sharing of resources for applications in materials modeling, exploiting supercomputing and cloud computing resources.⁵ The initial instantiation of Fenix is funded under a specific grant agreement of the HBP named the Interactive Computing E-Infrastructure (ICEI). Breakthrough scientific research exploiting Fenix resources has been documented at <https://fenix-ri.eu/infrastructures/successories>. For instance, an open source platform for constructing and simulating personalized brain network models using Fenix resources supports The Virtual Brain (TVB) workflows. The successful reconstruction and simulation of the cerebellar neurons and networks is another example. Both TVB and the Cerebellar Modelling Hub are part of a digital research infrastructure called EBRAINS, which has been included in the

2021 Roadmap of the European Strategy Forum on Research Infrastructures (ESFRI). The Fenix storage services have been used to share results of SARS-CoV-2 virus investigations. These projects highlight that Fenix resources not only enable access to HPC and cloud resources but are also on the critical path for realizing sustainable, digitalized research platforms for diverse scientific communities.

Federated Identity, Access, and Resource Management

Fenix leverages the European Authentication and Authorisation for Research and Collaboration (AARC) project's blueprint architecture for establishing federated identity and access management services.¹ As shown in Figure 2, the central proxy service is provided by GÉANT that manages one of the largest academic and research networks. The solution offers multiple levels of assurance and trust across the hosting sites as Identify Providers (IdPs) and communities such as

the HBP IdP. Fenix User and Resource Management Services (FURMS) provides federated access management to HPC and cloud resources. The core objectives of Fenix federation are a uniform experience for users, and extensibility such that Fenix AAI can be leveraged by the community or domain-specific platform development teams transparently. Fenix AAI facilitates identification and authentication of users by federating multiple IdPs, validating user profiles, maintaining a registry of usage agreement and policies, and the general Fenix usage agreement. FURMS provides central accounting, budgeting, and reporting mechanisms at different granularities (research groups or communities) and offers secure, role-based access controls. Furthermore, it serves as an attributes provider needed for authorization of Fenix services, for instance, secure ssh key management for HPC. 

References

1. AARC Community members, & AppInt members. AARC Blueprint Architecture 2019 (AARC-G045).

- Zenodo; <https://doi.org/10.5281/zenodo.3672785>.
2. Alam S. et al. Fenix: Distributed e-infrastructure services for EBRAINS. Brain-Inspired Computing. K. Amunts, L. Grandinetti, T. Lippert and N. Petkov N., Eds. *Lecture Notes in Computer Science 12339* (2021). Springer, Cham; https://doi.org/10.1007/978-3-030-82427-3_6
3. Alam S.R. et al. Archival data repository services to enable HPC and cloud workflows in a federated research e-infrastructure. In *Proceedings of the 2020 IEEE/ACM Intern. Workshop on Interoperability of Supercomputing and Cloud Tech.*
4. Amunts, K., Mohlberg, H., Bludau, S., Zilles, Julich-Brain: A 3D probabilistic atlas of the human brain's cytoarchitecture. *Science* (2020), New York, NY.
5. Talirz, L. et al. Materials Cloud, a platform for open computational science. *Sci Data* 7, 299 (2020); <https://doi.org/10.1038/s41597-020-00637-5>

Sadaf R. Alam is Chief Technology Officer of the Swiss National Supercomputing Centre (CSCS) at ETH Zürich, Switzerland.


Javier Bartolome is Systems Group Manager at the Barcelona Supercomputing Center (BSC), Spain.

Michele Carpenè is Software Developer and Project Manager at the Italian Supercomputing Centre (CINECA), Italy.

Kalle Happonen is Development Manager at the Finnish Supercomputing Center (CSC), Espoo, Finland.

Jacques-Charles Lafoucrière is Program Manager at the Commissariat à l'Énergie atomique et aux énergies alternatives (CEA), France.

Dirk Pleiter is a professor of computer science at KTH and Director of PDC, Sweden and Senior Researcher, Juelich Supercomputing Centre (JSC), Germany.

 This work is licensed under a Creative Commons Attribution 4.0 International License. <http://creativecommons.org/licenses/by/4.0/>

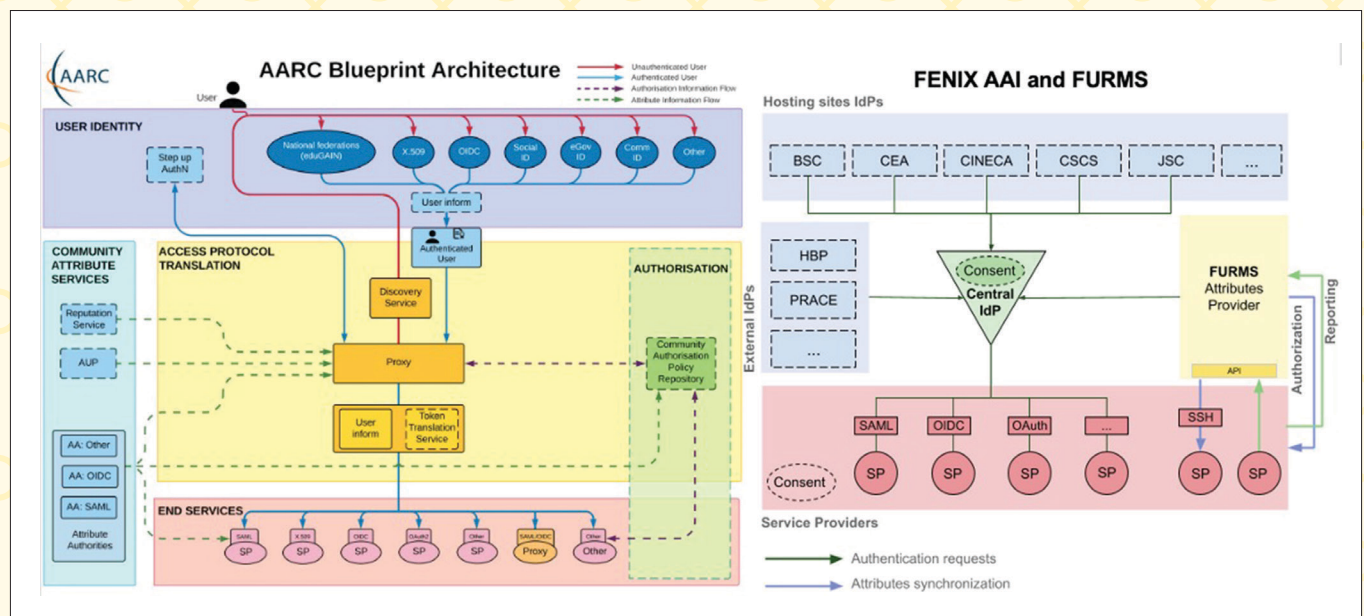


Figure 2. Fenix AAI and FURMS implementation compliant with AARC Blueprint Architecture.