



FENIX

RESEARCH INFRASTRUCTURE

Archival Data Repository Services to Enable HPC and Cloud Workflows in a Federated Research e-Infrastructure

Sadaf Alam (CSCS), Dirk Pleiter (JSC) and others



The ICEI project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 800858.

Outline

- Background on Fenix
- Fenix compute and data services
- Approach to federation
- Selected case study
- Status and outlook

Background: Human Brain Project




- Flagship project funded by European Commission with the goal of realising a infrastructure for brain-related research

- Neuroscience, brain-inspired computing, brain-related medicine

- New research infrastructure: EBRAINS

- Provides an increasing set of services in the area of data&knowledge, atlases, simulation

- Key challenge: Diverse community with vastly different needs



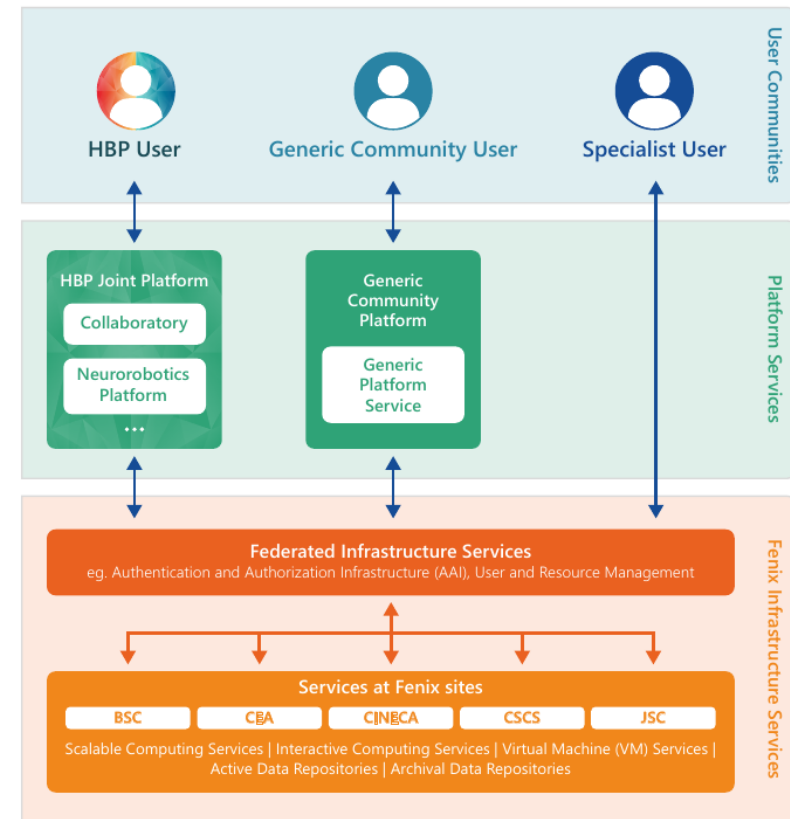
[CEA/INRIA, 2020]



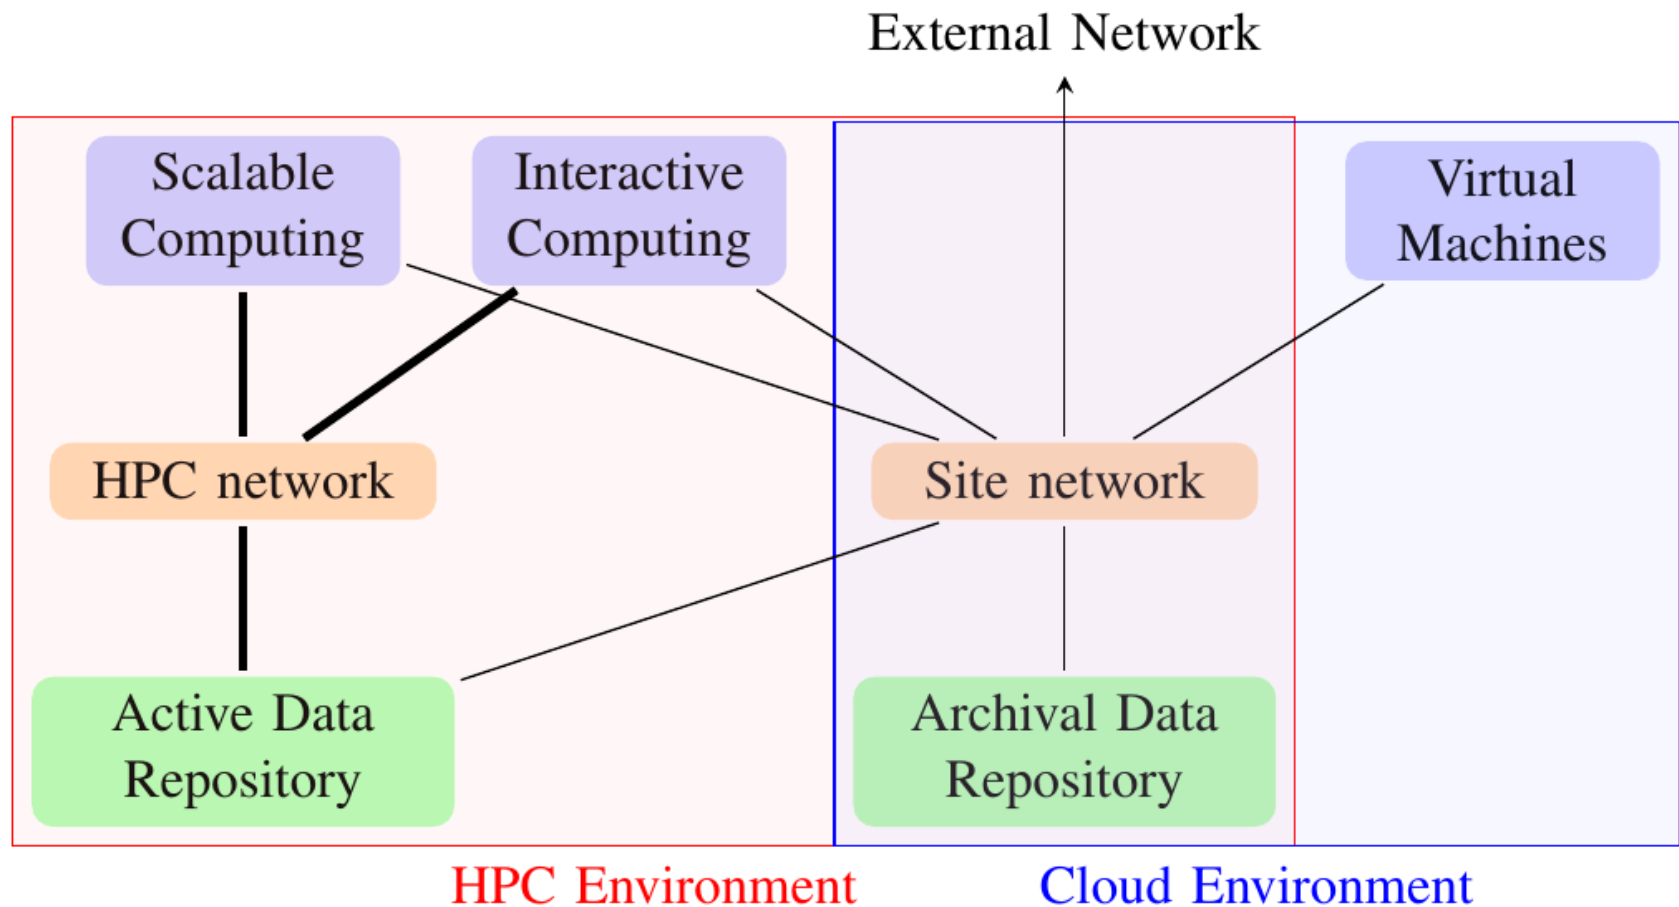
EBRAINS

Background: Fenix

- Effort to harmonise and federate service offerings between supercomputing centres
 - Currently 5 supercomputing centres involved
- Serving multiple user communities
 - HBP is the current prime community
- Concept of layered organisation of services
 - Infrastructure versus platform services layer



Fenix Compute and Data Services



Fenix Compute Services

■ Scalable Computing Services

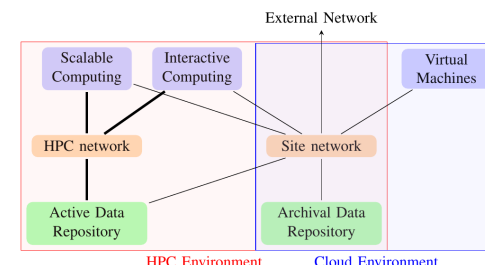
- Massively parallel HPC systems for scalable and/or compute heavy applications
- Examples: Simulations, ML/DL workloads

■ Interactive Computing Services

- Interactive access to fast servers and large-scale data-sets
- Support of interactive frameworks like Jupyter notebooks

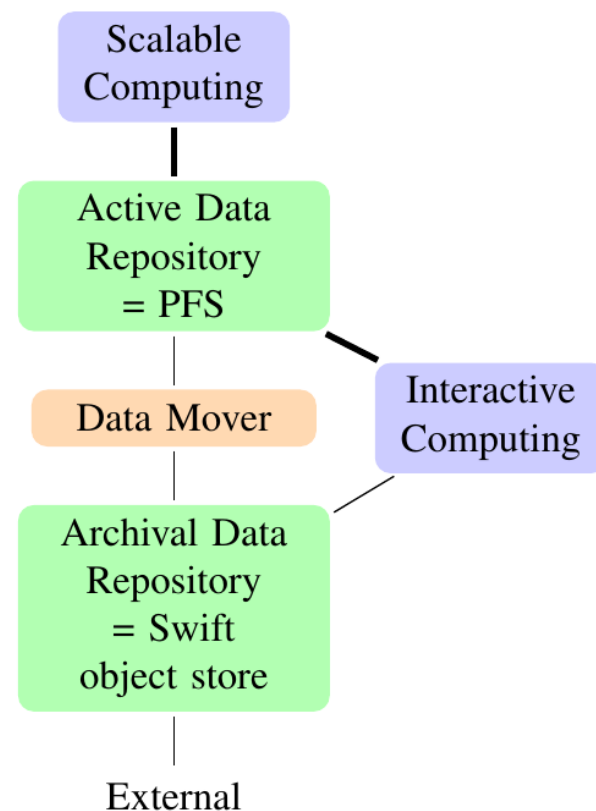
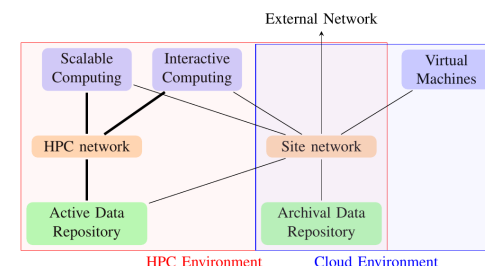
■ Virtual Machine Services

- Service for deploying virtual machines in a stable and controlled environment
- Example: Deployment of web services for data sharing



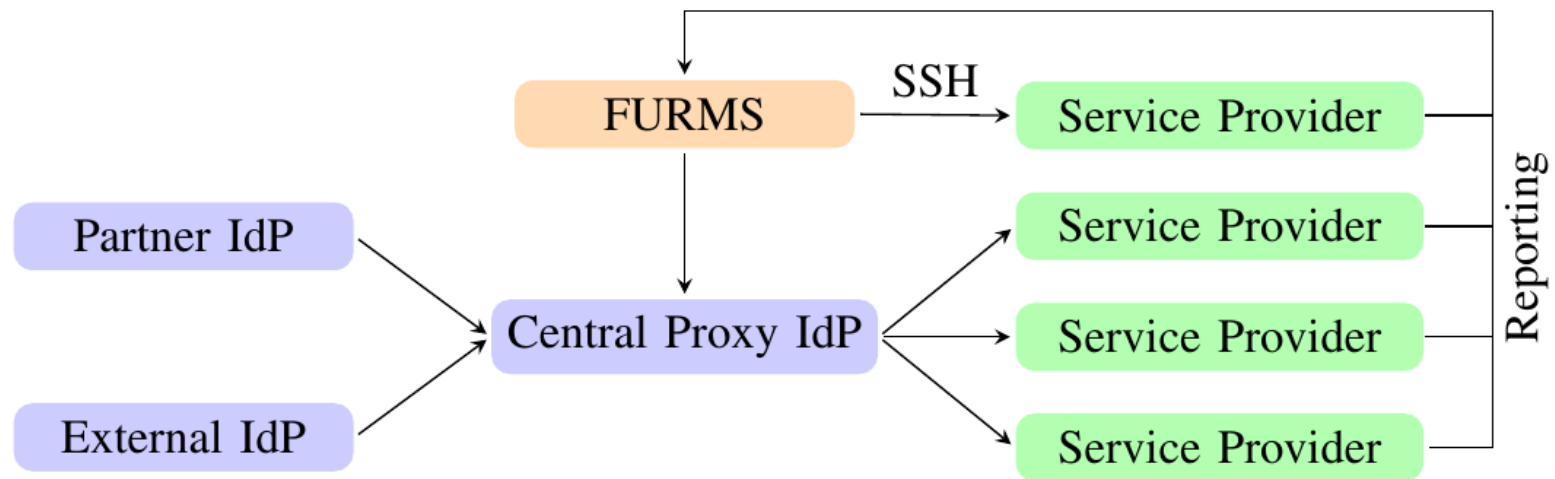
Fenix Data Services

- Federated Archival Data Repositories (ARD) with Cloud-standard access interfaces
 - Suitable for large-scale data sets
 - Access via Swift interface
- Non-federated Active Data Repositories (ACD) optimised for performance
 - Typically standard parallel file system
- Data mover to move data asynchronously between ARD and ACD
 - Site-local Swift-POSIX copy



Approach to Federation: AAI (1/2)

- Federation of IdPs via Central Proxy IdP
 - Support for SAML and OIDC
- Fenix User and Resource Management Service (FURMS)
 - Provider of SSH keys



Approach to Federation: AAI (2/2)

- Federation via Central Proxy IdP
 - Federation of multiple IdPs by proxying authentication requests
 - Validation of user profiles
 - Policy registry and management of principles of engagement
 - Managing of general Fenix Usage Agreement
- Multiple levels of assurance

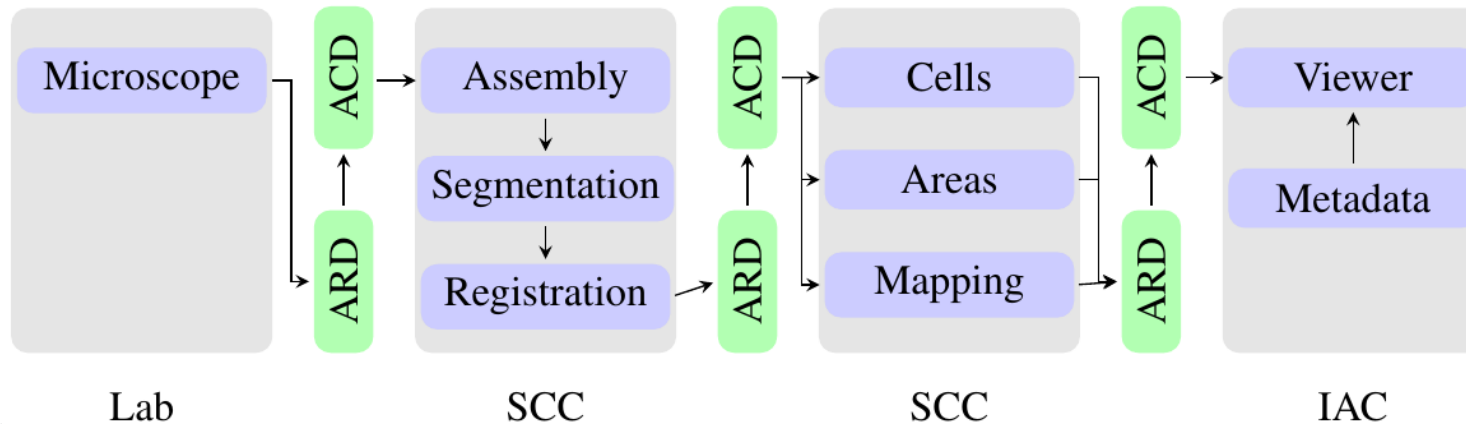
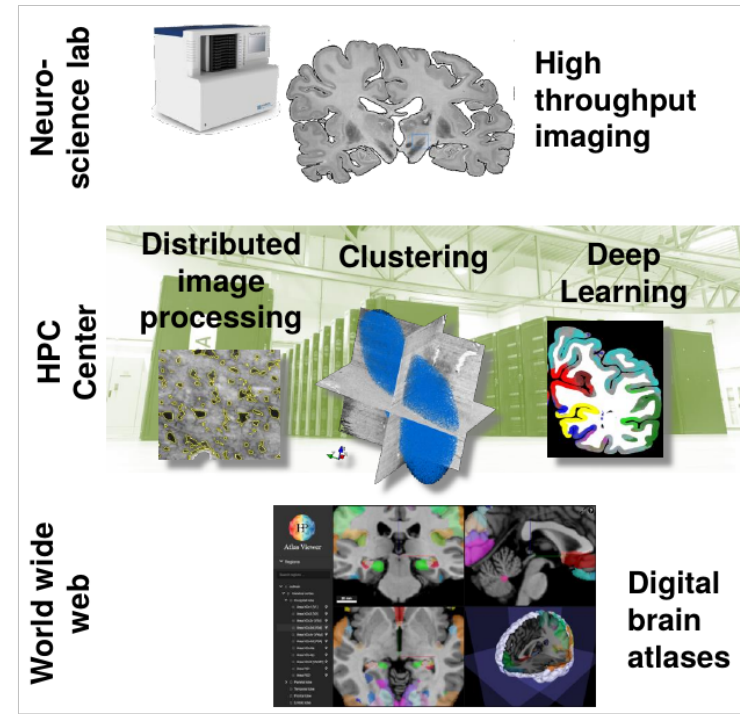
HIGH	HIGH (IdPs of supercomputing centres)
MEDIUM	MEDIUM (community IdPs, e.g. Human Brain, PRACE)
LOW	other

Approach to Federation: FURMS

- Generic attribute provider
 - Policy document management
 - SSH key management
- Membership management service
 - Support of different user roles
- Resource allocation to communities and projects + accounting
 - Fenix credits with the following attributes
 - Type of resource
 - Amount of this resource
 - Fenix resource provider
 - Validity period

Case Study: Brain Atlas

[T. Dickscheid]



Status

- Resources available for research projects
 - Equipment deployed at CSCS, JSC, CEA
 - Upcoming resources at BSC, CINECA
- Resource allocation mechanisms in place
 - EBRAINS and PRACE-ICEI allocation process
- Initial version of AAI services in place
- Different components still under development
 - Data Mover Services, FURMS

Summary and Outlook

- Fenix as an effort for integration of HPC and Cloud services
 - Federation of services of different supercomputing centres
- Currents efforts focus primarily on brain research community through the HBP
 - Generalised concept for supporting other communities
- Future steps
 - Expand number of involved sites
 - Align with efforts in the context of European Open Science Cloud



Thank you!