FENIX RI

FET FLAGSHIPS

Human Brain Project

Co-funded by the European Union

Federated HPC, Cloud and storage services for research in Europe

Dirk Pleiter (Forschungszentrum Jülich/JSC)

Heidelberg 25.11.2019

Fenix Concept

- Service multiple **user communities** (including the brain research community)
- Provide **federated infrastructure services** to support **community specific platform services** (e.g. EBRAINS services)
- **Site local infrastructure services** versus **federation services**
- **Composability of services** according to the needs of users
- Support for **different local implementations** of services to avoid technology/vendor lock

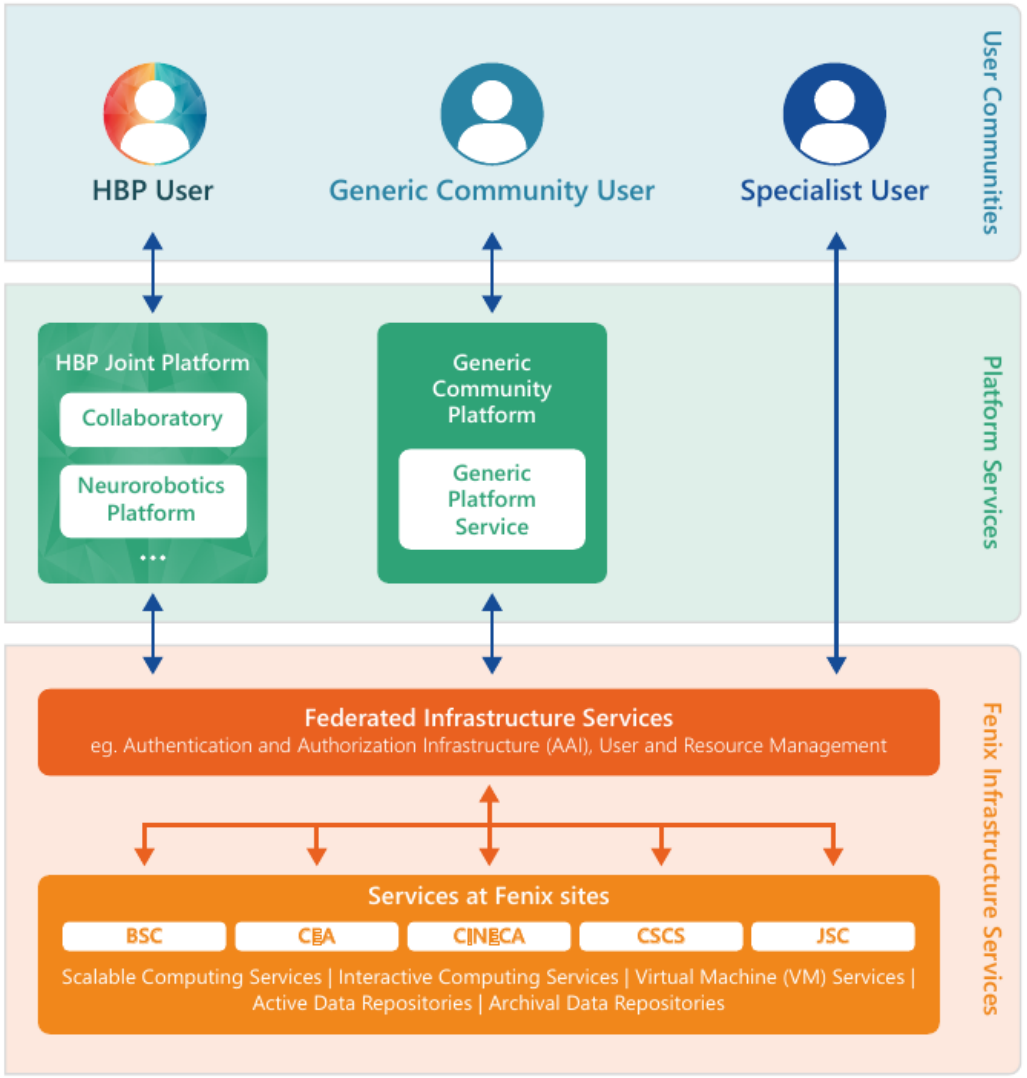

Co-funded by the European Union

Key Fenix Services

- **Computing services**
  - Interactive Computing Services
  - Scalable Computing Services
  - Virtual Machine Services
- **Data services**
  - Active Data Repositories
  - (Federated) Archival Data Repositories
  - Data Mover Services, Data Location and Transport Services
- **Federation services**
  - Authentication and Authorisation Services (AAI)
  - User and Resource Management Services (FURMS)

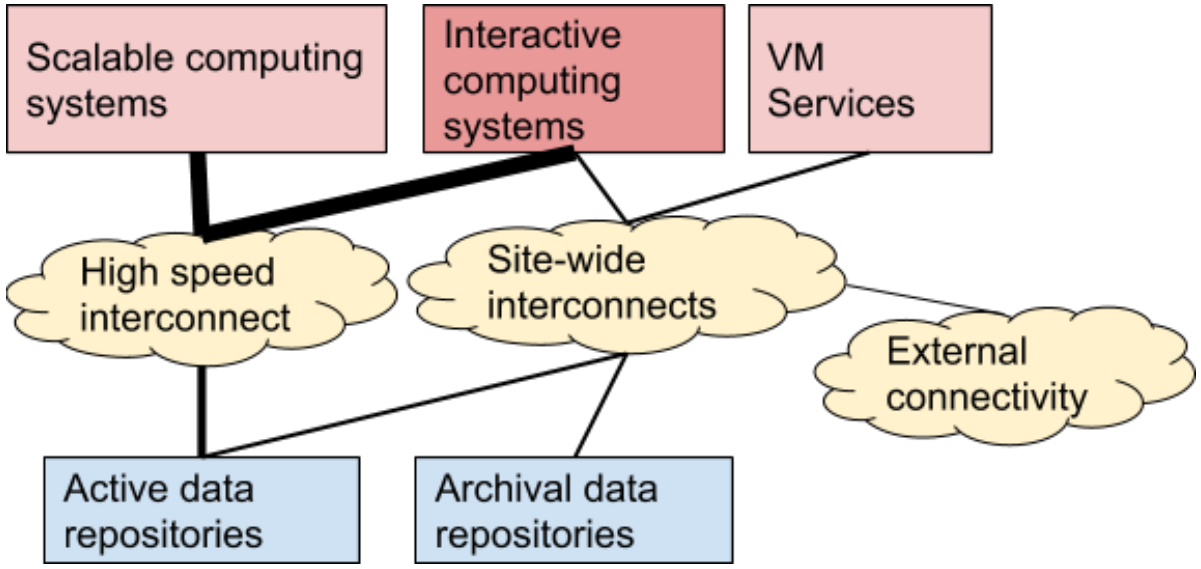

Co-funded by the European Union

Fenix Resource Providers

- Partners providing resources today/soon
  - Barcelona Supercomputing Centre (Barcelona, ES)
  - CEA (Bruyeres-le-Chatel, FR)
  - CINECA (Bologna, IT)
  - CSCS (Lugano, CH)
  - Jülich Supercomputing Centre (Jülich, DE)
- Extensible to other sites
  - Not limited to supercomputing centres

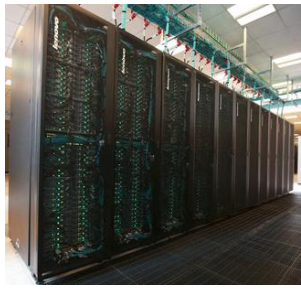

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CINECA

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

JÜLICH
Forschungszentrum

Co-funded by the European Union

Benefits of a Federated e-Infrastructure

- Data locality
  - Facilitate processing of data where data is located
- Higher availability of services
  - Services can be replicated at different sites
- Diversity of infrastructure service implementations
  - Local resource providers provide different service implementations
- Diversity of expertise
  - Different kind of experts at different resource provider sites
- Improved sustainability of overall infrastructure
  - Individual resource providers are (partially) replaceable

Compute Services

- Scalable Computing Services
  - Massively parallel HPC systems for scalable applications
  - Possibly elastic access
- 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
  - Fully customisable by user (e.g. platform developer)

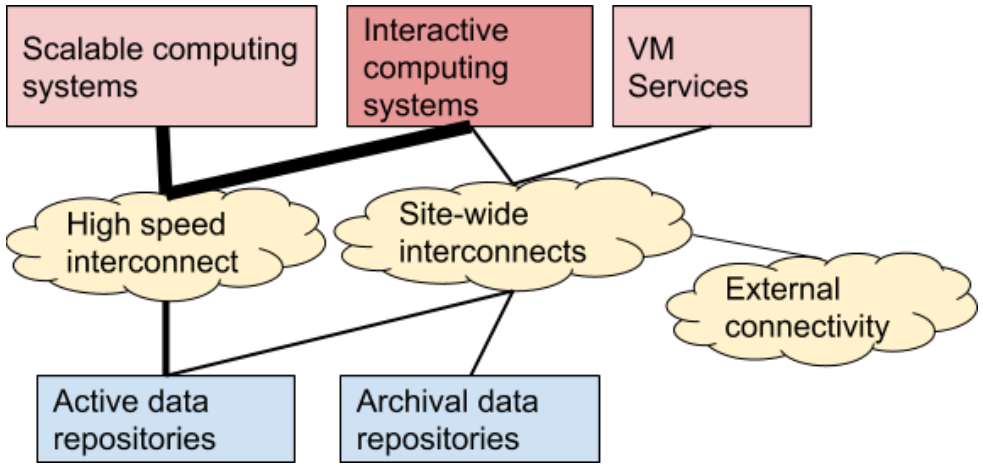

Data Services

- Federate **Archival Data Repositories** with Cloud-standard access interfaces
  - Access via SWIFT interface
- Non-federated **Active Data Repositories** optimised for performance
  - Typically POSIX interface
- **Data mover** to move data asynchronously between data repositories
- Support for storing sensitive data + long-term archiving of data

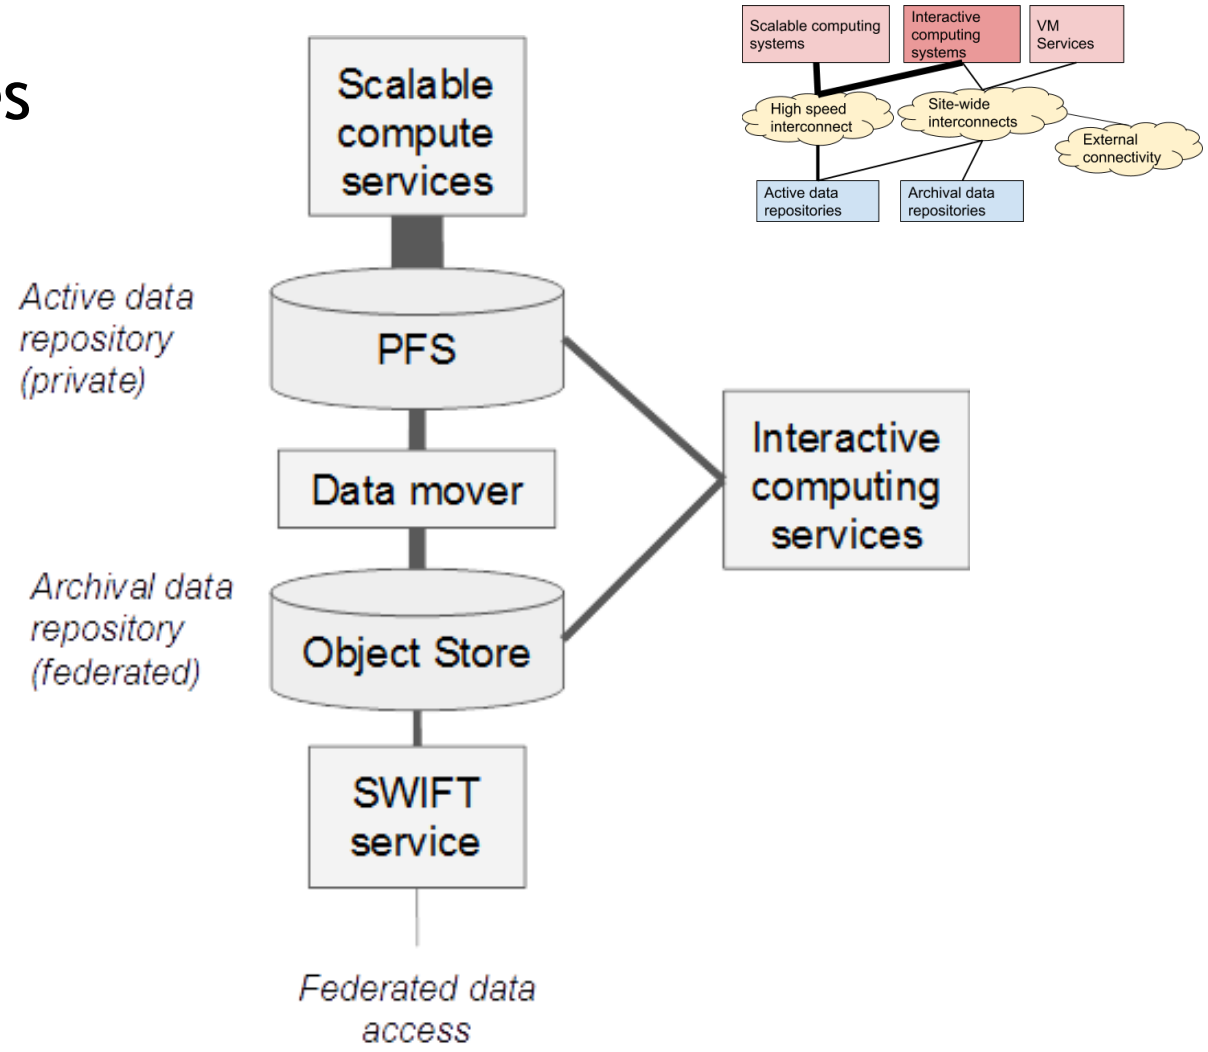

Co-funded by the European Union

Federation Services

- Authentication and Authorisation Services
  - Support industry standard protocol (e.g. OIDC, SAML2)
  - Central proxy Identity Provider (IdP)
  - Trust the IdPs of all resource providers
    - Open for other IdPs
- Fenix User and Resource Management Service (FURMS)
  - Group/project membership management
  - Authorization attributes provider
  - SSH public keys management
  - Managing site specific Usage Agreements
  - Reporting and metering

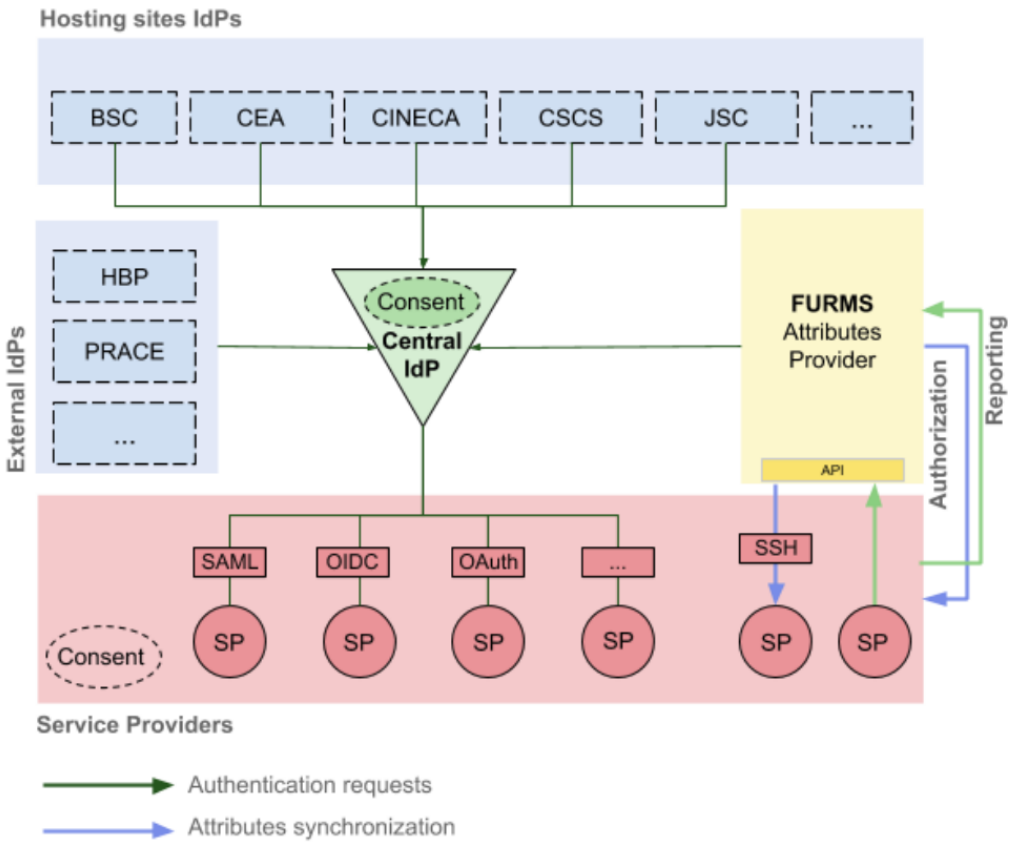

Co-funded by the European Union

Implementation at CSCS: Overview

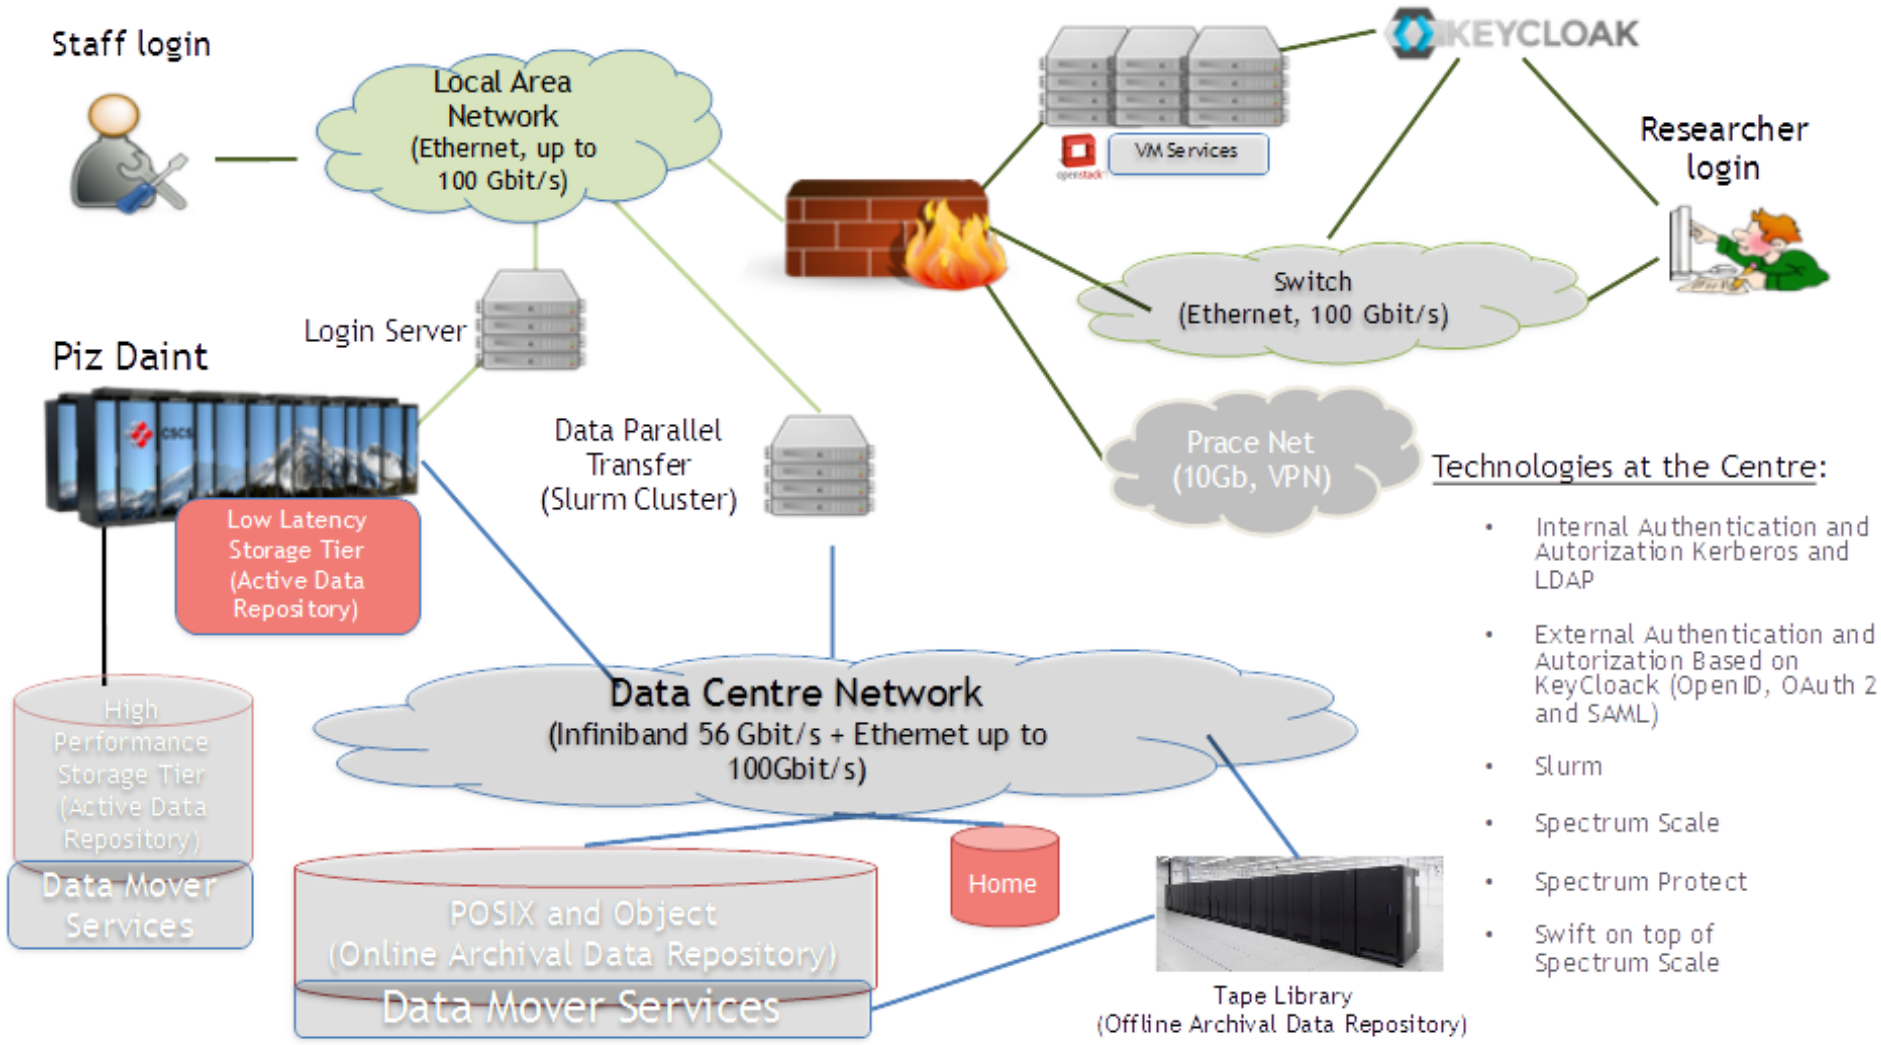

Co-funded by the European Union

Implementation at CSCS: Details

- Piz Daint Multicore
  - 2x Intel Xeon E5-2695 v4 (2.10GHz, 18 cores)
  - 64 or 128 GByte host memory
  - Up to 1431 nodes per job

- Piz Daint Hybrid
  - 1x Intel Xeon E5-2690 v3 (2.6 GHz, 12 cores) + 1x NVIDIA P100 GPU
  - 64 GByte host + 16 GByte device memory
  - Up to 5320 nodes per job
- Active Data Repositories
  - Lustre parallel file system
- Archival Data Repository
  - Access via Swift interface

Co-funded by the European Union

Fenix Resource Allocation (1/2)

- **Fenix Credits**

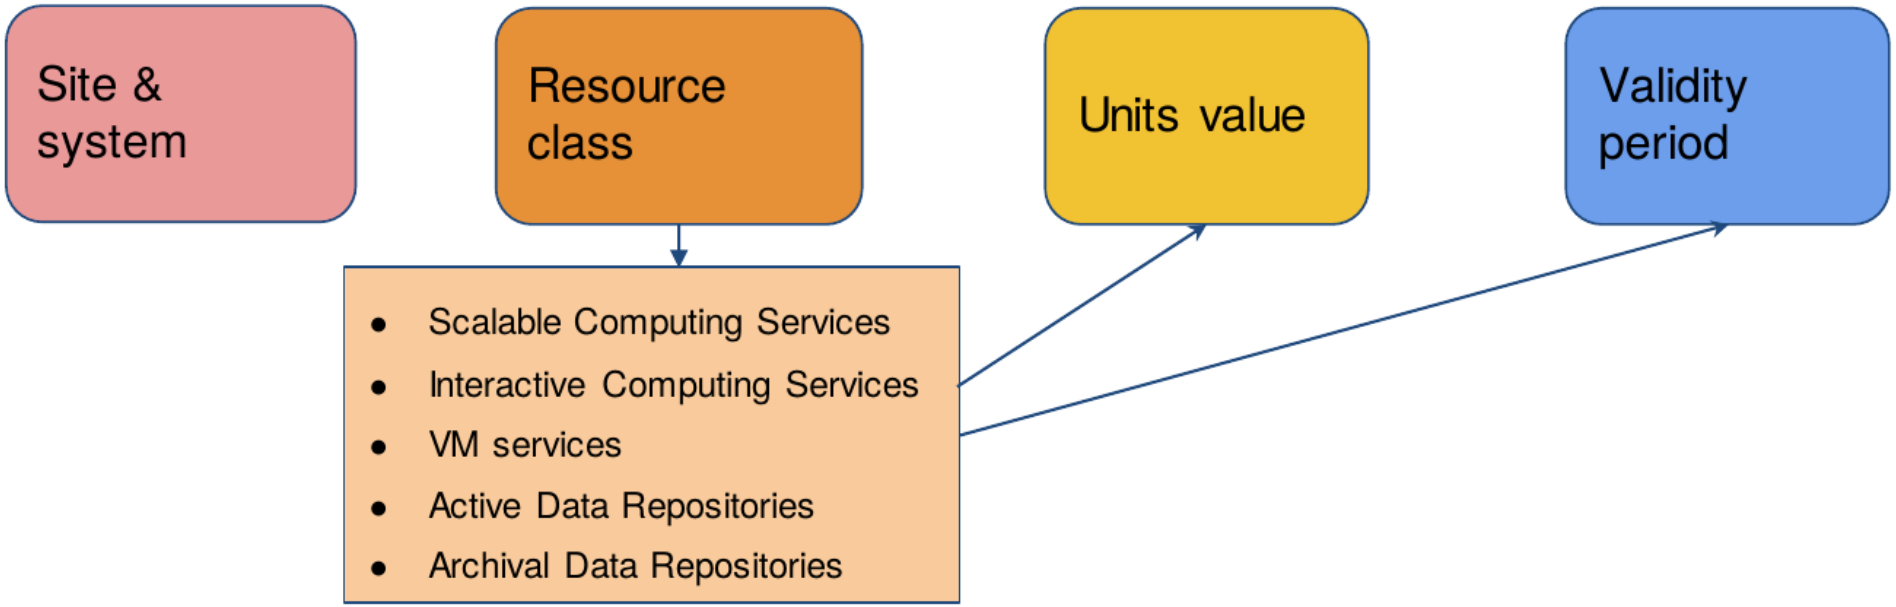

- Fenix Credits allocated to **Fenix Communities**
- Fenix Communities allocate Fenix Credits to **Fenix Projects**

Co-funded by the European Union

Fenix Resource Allocation (2/2)

- Principles of the allocation mechanism for resources that are available through Fenix:
  - Process follows peer review principles established by PRACE
  - Each Fenix Community (e.g. HBP) is responsible for the actual distribution of their share within that community
- HBP is the initial prime and lead user community
  - 25% of available resources are reserved for HBP
  - 15% are provided to European researchers at large via PRACE
  - The remaining 60% are with the respective centre that is providing the resources and are made available to users, e.g. via national calls

- Access for HBP members
  - Continues call for proposals with decisions on resource allocations managed by EBRAINS Infrastructure Allocation Committee
- Access for non-HBP members
  - ICEI PRACE Tier-1 calls opened 4x per year

Co-funded by the European Union

Usage Examples/Scenarios (1/2)

- Large scale simulations of models: Hippocampus
  - Project studies the mechanisms that may contribute to the emergence of higher brain functions at the cellular and behavioural level in the hippocampus
- Required resources of the Fenix Infrastructure
  - **Scalable Computing Services** for running large-scale simulations using Neuron
  - **Active Data Repositories** as temporary storage to write from simulation and read for analysis
  - **Interactive Computing Services** for analysing data produced by simulations
  - **Archival Data Repositories** for storing final data products
  - **VM Services** for providing web-based access to visualisations of the final data products

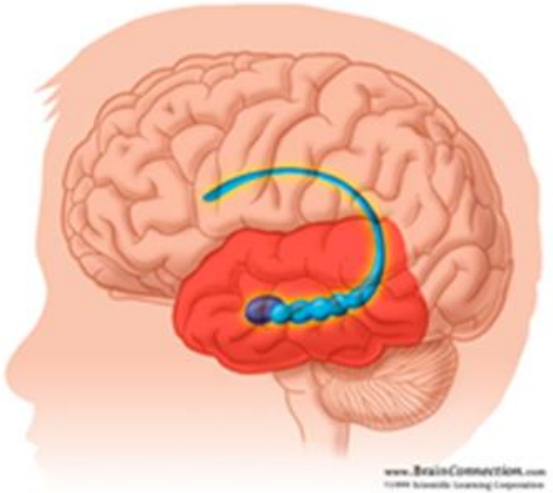

[Posit Science, 2010]

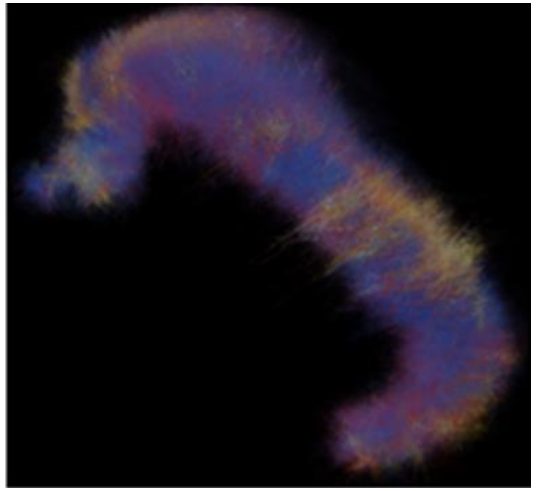

[M. Migliore et al.]

Usage Examples/Scenarios (2/2)

- Creation of a digital brain atlas
  - Input: O(1) PByte of brain slice image data
- Required resources of the Fenix Infrastructure
  - Archival Data Repositories for storing image data
  - Interactive Computing Services for interactive analysis of images
  - Scalable Computing Services, e.g., for
    - Image segmentation based on deep learning methods
    - Image registration
  - Active Data Repositories as temporary storage for fast access to data
  - VM Services for providing web-based access to data

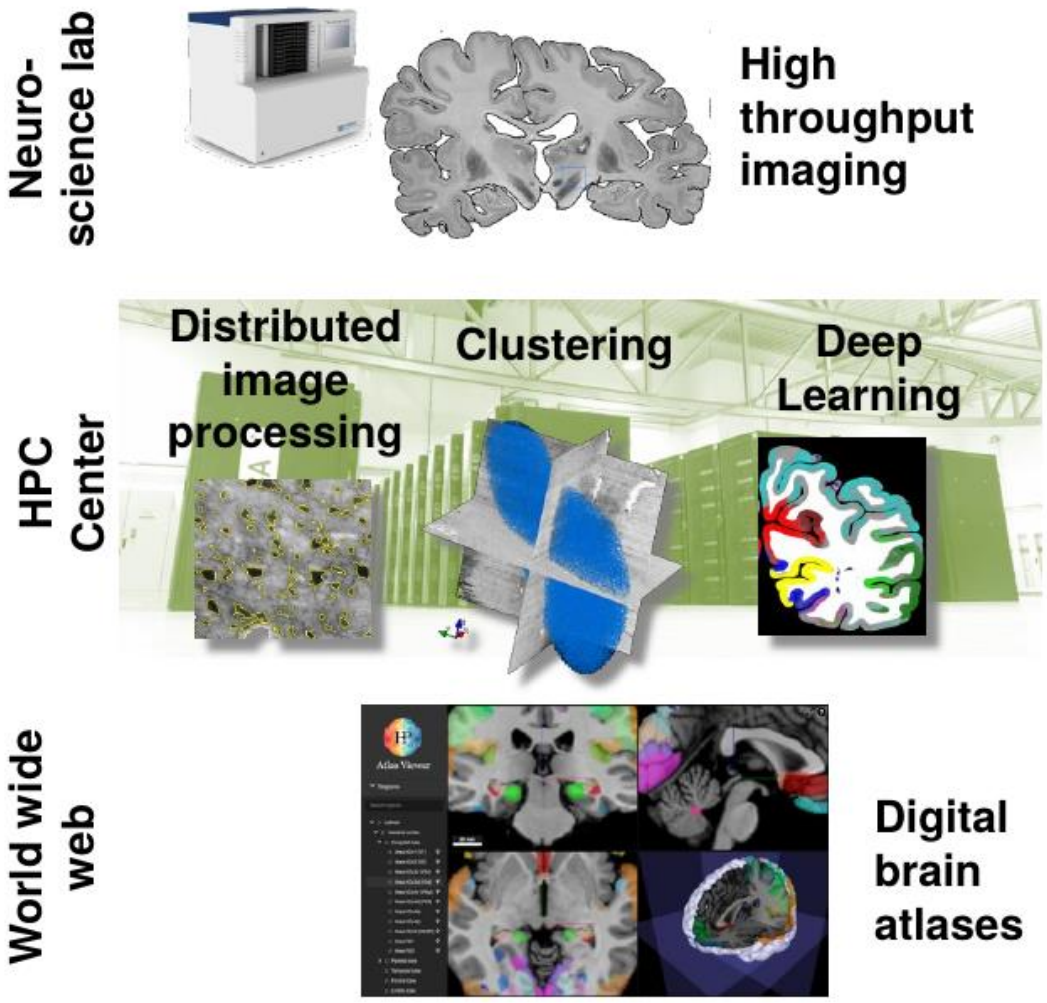

[T. Dickscheid]

Summary and Outlook

- Fenix provides a generic concept for federated e-infrastructure services
  - Basic layer for deploying community specific platform services like EBRAINS services
- Initial set of Fenix services available for brain researchers and European researchers at large
  - Allocation of resources through HBP and PRACE
- Service portfolio is being extended to 5 European supercomputing sites
  - Expect in future other resource providers to join

Human Brain Project

Co-funded by
the European Union

Thank You

www.humanbrainproject.eu

@HumanBrainProj

Human Brain Project