



# FENIX

RESEARCH INFRASTRUCTURE

## Introduction to Fenix and ICEI

Dirk Pleiter, JUELICH and Colin McMurtrie, CSCS/ETH Zurich

Sophia, Bulgaria

14 May 2018



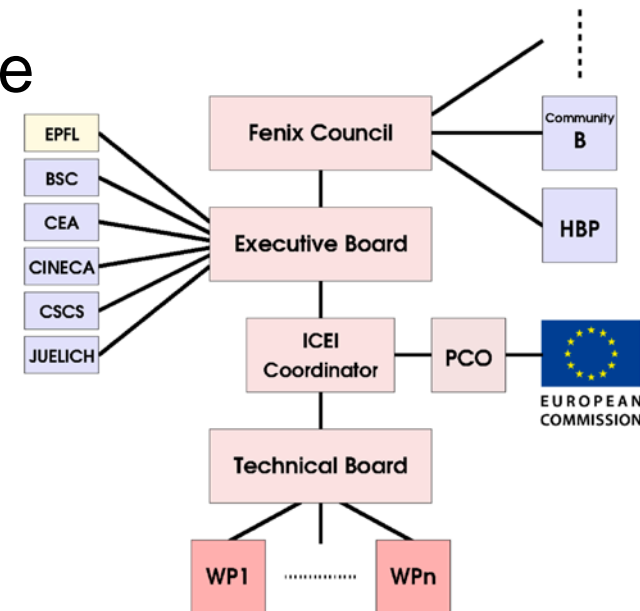
The ICEI project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 800858.

# Fenix Goals

- Establish HPC and data **infrastructure services** for multiple research communities
  - Encourage communities to build community specific platforms
  - Delegate resource allocation to communities
- Develop and deploy services that facilitate **federation**
  - Based on European and national resources
- **Science community driven approach**
  - Infrastructure realisation and enhancements based on co-design approach
  - Science communities providing resources to realise infrastructure  
→ HBP SGA Interactive Computing E-Infrastructure
  - Resource allocation managed by community

# Fenix Governance Structure

- Fenix is a Consortium
  - Establishing an **MoU** between the participating HPC sites
    - Comprised of large HPC sites within Europe
    - Currently, BSC, CEA, CINECA, CSCS, JSC
    - The intention is to keep the number of participating sites below 10
  - EPFL included for the ICEI part only
    - Fenix has close ties to the HBP via ICEI (next slide)



# ICEI as the First Instantiation of Fenix

- The Human Brain Project (HBP) presented a unique opportunity
  - Funding available from the EC for a **federated data and compute infrastructure**
  - Name: **Interactive Computing e-Infrastructure (ICEI)**
  - Nice fit with the concept for Fenix which had grown out of the **Federated Data Pilot Project (FeDaPP)** run in 2016 at the start of HBP SGA1
- ICEI is the **first Instantiation** of the Fenix Infrastructure
  - Funds the development of many of the federated services in Fenix
  - HBP-centric but Fenix has the goal to be more general

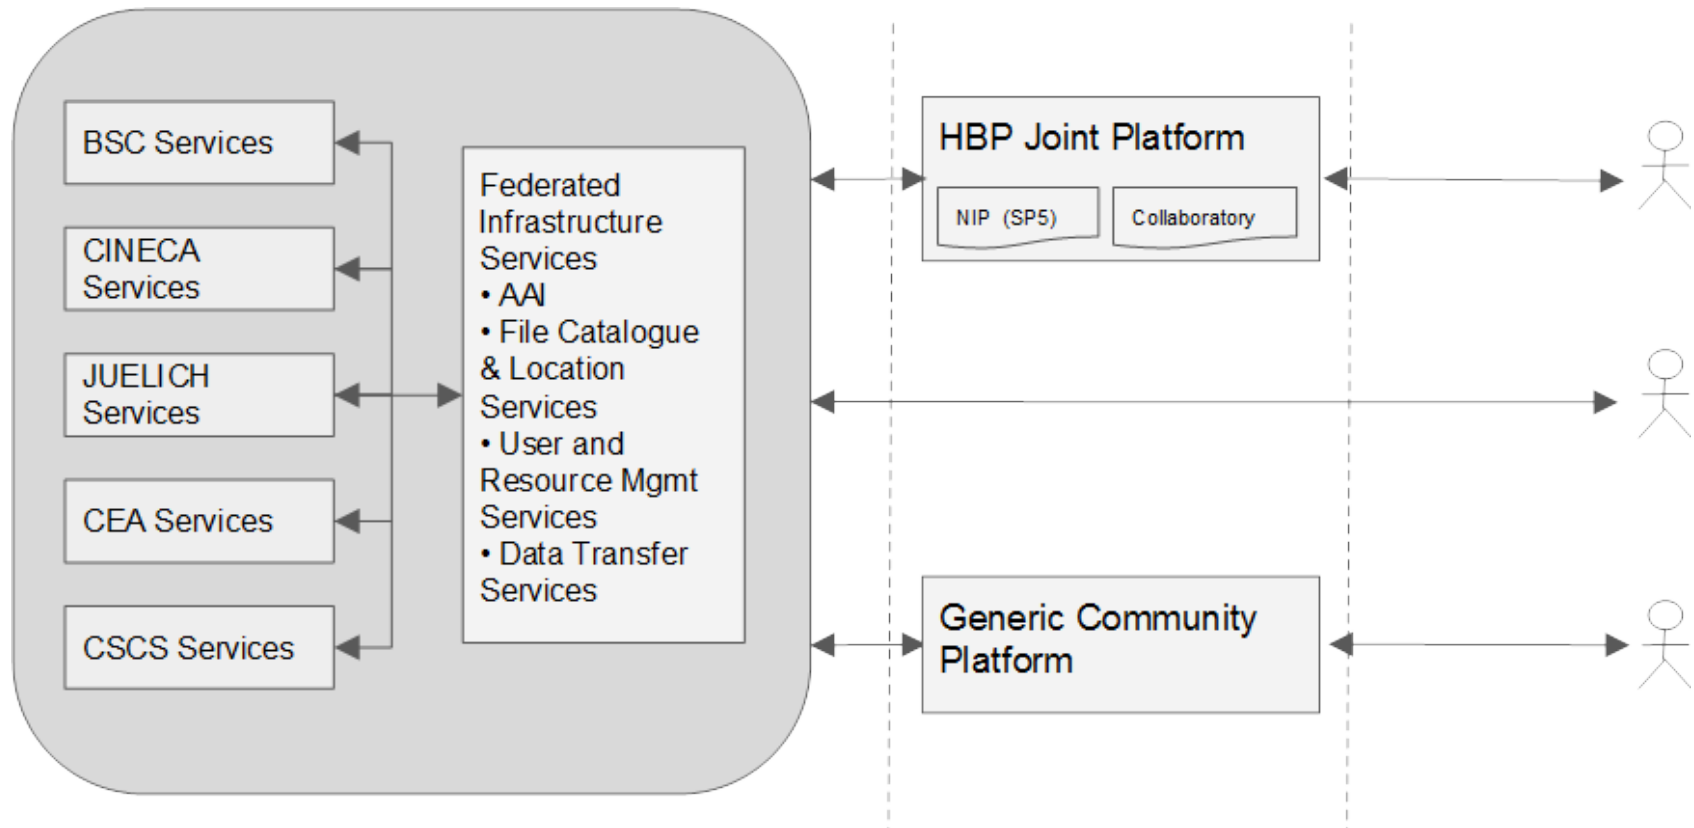
# ICEI Summary

- ICEI is an SGA under the HBP FPA-CA
  - EPFL is the coordinator of HBP and hence plays a role in ICEI coordination
- ICEI to provide **infrastructure services** to HBP
  - HBP SP7 (High-performance Analytics and Computing Platform) is in the role of developer/provider of platform services
  - Strong links between ICEI and SP7 DoA
- Timeline overview
  - Evaluation of proposal finalised by the EC in 4Q2017
  - Project duration of **63 months** (1.1.2018 – 31.3.2023)
  - First resources available at CSCS as of April 2018
  - First **co-design workshop** took place on 9.2.2018
  - **Public Information event** held 15 March 2018
  - **RFI meetings** held with vendors 26-27 May 2018

# Fenix Architectural Concept (1/2)

- Service-oriented provisioning of resources
  - Focus on infrastructure services suitable for different science communities
- Support for community specific platforms
  - Encourage and facilitate community efforts
- Federation of infrastructure services
  - Enhance availability of infrastructure services
  - Broaden variety of available services
  - Optimise for data locality

# Architectural Concept (2/2)



**ICEI infrastructure services**

**SP7 platform services**

# Overview over Planned Fenix Services

- Computing services
  - Interactive Computing Services
  - (Elastic) Scalable Computing Services
  - VM Services
- Data services
  - Active Data Repositories + (federated) Archival Data Repositories
  - Data Mover Services, Data Location and Transport Services
- Other
  - Authentication and Authorisation Services
  - User and Resource Management Services (FURMS)
  - Monitoring Services
  - Internal/external interconnect



# Interactive Computing Services

- Interactivity
  - Interactive processing of data
  - Capability of a system to support distributed computing workloads while permitting
    - Monitoring of applications
    - On-the-fly interruption/steering by the user
- Architectural requirements
  - Interactive access
  - Tight integration with scalable compute resources
  - Fast access to storage resources
- Support for interactive user frameworks
  - Jupyter notebook, R, Matlab/Octave



# (Elastic) Scalable Computing Services

- Different options for service provisioning
  - Access to highly scalable compute resources with possible longer wait times
  - Elastic access to a limited amount of compute resources
- Possible realisation of elastic provisioning
  - Free resources by means of checkpoint/resume mechanisms
  - Reserve (small) amount of nodes
- Considered use case
  - Coupling of neuro-robotics experiments to brain simulations
- Open co-design questions
  - Upper limit for acceptable response times
  - Scaling range

# Virtual Machine Services

## ■ Use case

- Deployment of **community platform services** running 24/7
- Examples: HBP *Collaboratory*, database servers, visualisation services

## ■ Requirements

- Allow communities to flexibly create and manage VM services similar to a cloud environment
- Provide stable infrastructure services

# Architectural Concepts: Data Store Types

## ■ Archival Data Repository

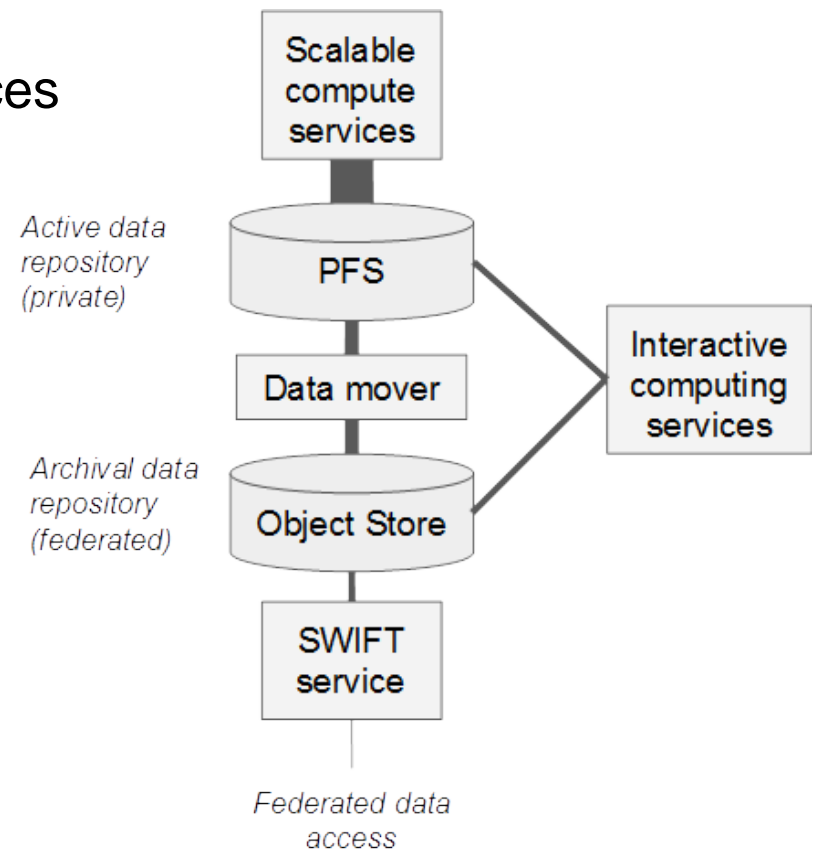
- Data store optimised for **capacity, reliability and availability**
- Used for storing large data products permanently that cannot be easily regenerated

## ■ Active Data Repository

- Data repository localized close to computational or visualization resources **optimised for performance**
- Used for storing temporary slave replica of large data objects

# Storage Architecture

- Concept
  - **Federated archival-class data repositories** with Cloud interfaces
  - **Non-federated active data repositories** with POSIX interface accessible from HPC nodes
- Envisaged implementation
  - Mandate same technology at all sites
  - Current prime candidate: **OpenStack SWIFT**



# Resource Allocation Model

- Actors
  - Fenix Resource Providers
  - Fenix Communities
  - Fenix Users
- Role of Fenix Resource Providers
  - Provide fixed amount of resources for given period to Fenix Communities
  - Define rules for resource allocation (e.g., peer-review process)
- Fenix Users
  - Submit proposal for resources to relevant Fenix Community
- Fenix Community
  - Review proposal and award available resources to Fenix Users

# Fenix Credits

- Fenix Credit =
  - Vouchers for authorising resource consumption
- Different types of resources
  - Scalable compute resources ( $N_{\text{node}} \times \text{time}$ )
  - Interactive computing services ( $N_{\text{node}} \times \text{time}$ )
  - Active data repositories (capacity  $\times$  time)
  - Archival data repositories (capacity)
  - Virtual Machines (VM flavours[cores, memory, network IPs]  $\times$  time)
- Credit attributes
  - Value and type of resource
  - Fenix Resource Provider
  - Validity period

## Note

**Fenix Credits** will only allow to use the ICEI infrastructure. Other access mechanisms apply for other resources (e.g. PRACE)

# User Management

## ■ Model

- Scientist identifies itself through virtual identity issued by accepted Identity Provider
- Scientist registers with Fenix Community to become a Fenix User

## ■ Workflow

- Scientist obtains virtual identity
- Scientist applies for membership in a Fenix Community and accepts Fenix Community Usage Agreement
- Fenix Community decides on application





# FENIX

RESEARCH INFRASTRUCTURE

Q&A

[colin.mcmurtrie@cscs.ch](mailto:colin.mcmurtrie@cscs.ch)

Fenix Website:

<https://fenix-ri.eu>