FENIX

RESEARCH INFRASTRUCTURE

HPC and Cloud Infrastructure for Neuroscientists

Webinar

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

Outline

- The ICEI Project
- Fenix Infrastructure
  - Overview
  - Services
  - Available resources
  - User access
  - Outlook
- Use case example(s)
- Getting support
- Q&A

FENIX RI

The ICEI Project

- Interactive Computing E-Infrastructure for the Human Brain Project (HBP)
  - Co-funded by the European Commission, Specific Grant Agreement under the umbrella of the HBP Framework Partnership Agreement
  - Started officially in January 2018
  - 5 Supercomputing centres (PRACE Hosting Members) involved:

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

DE LA RECHERCHE À L'INDUSTRIE
cea

CINECA

ETH zürich

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

JÜLICH
Forschungszentrum

  - The **ICEI project** carries out coordinated procurements of equipment and R&D services to realise elements of the **Fenix Infrastructure**

FENIX RI

Fenix Infrastructure – Overview

- **Federated computing and data services** for European researchers
- **Data storage and scalable computing resources** in **close proximity** to each other and tightly integrated
- **Service-oriented** provisioning of resources, aiming to
  - Meet the requirements of various science communities
  - Form a basis for the development and operation of community-specific platform tools and services
- **Federation** of infrastructure services to **optimize for data locality**, **enhance availability** and **broaden variety** of services

FENIX RI

Fenix Infrastructure – Services

- **Scalable Computing** Services (SCC)
  - Massively parallel HPC systems suitable for large-scale brain simulations and high-throughput data analysis tasks
- **Interactive Computing** Services (IAC)
  - Quick access to single compute servers to analyse and visualise data interactively, or to connect to running simulations using SCC
- **Virtual Machine (VM)** Services
  - Service for deploying VMs in a stable and controlled environment, e.g. platform services like the HBP Collaboratory
- **Active Data Repositories** (ACD)
  - Site-local data repositories for storing temporary slave replicas of large data sets (parallel file systems)
- **Archival Data Repositories** (ARD)
  - Federated data store for long-term storage and sharing of large data sets

Fenix Infrastructure – Available resources

- Fenix Infrastructure services are available only at **CSCS**, yet. Overview of **available resources**:

Component	Service Type	ICEI Total Allocation (100%)	Available Resources (quarterly)
Piz Daint Multicore	SCC	250 nodes	465’375 node-hrs
Piz Daint Hybrid	SCC + IAC	400 nodes	744’600 node-hrs
OpenStack IaaS	VM	35 servers	35 servers
POSIX, Object and Tape	ARD	4 PB	4 PB
Low-Latency Storage Tier	ACD	80 TB	80 TB

- **Resources at other centres** will become available within the next months, details are published on the Fenix website: https://fenix-ri.eu/infrastructure/resources/planned-resources

Resources at CSCS – Some 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
  - 2x 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

- Archival Data Repository
  - Access via Swift interface
  - Same interface at all sites

Fenix Infrastructure – User access

- **Allocation Mechanism** for resources that are available within the Fenix Infrastructure, principles:
  - Process follows **peer review principles** established by **PRACE**
  - Each **user 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
  - Flyer "**Access to Fenix IT Services for HBP Users**" on Fenix website
- **Access for non-HBP** members
  - **PRACE Tier-0** calls and in the future dedicated **ICEI PRACE Tier-1** calls

Fenix Infrastructure – Outlook

- **Infrastructure components** that are procured within ICEI are becoming **operational at all centres**
- Realisation of **federated infrastructure services**, including:
  - **Authentication and Authorization Infrastructure** (AAI)
    - Enabling central user identification and authentication
    - Fenix AAI will provide access to infrastructure services, e.g. storage, while community platforms and services will retain their own authentication and authorization mechanisms
  - **Fenix User and Resource Managements Service** (FURMS)
    - Central system to manage membership and roles of users within certain stakeholder groups
    - Allows for resource allocation and accounting within these groups
  - **Data location and transfer** services
    - Enabling users of the Fenix infrastructure to locate and **move their data between centres**
  - **Data Mover** Service
    - Enabling users to **move data between ACDs and ARDs** at a **single centre**
- Infrastructure is open to **new user communities and centres**

Use case example

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 *(PI: M. Migliore)*

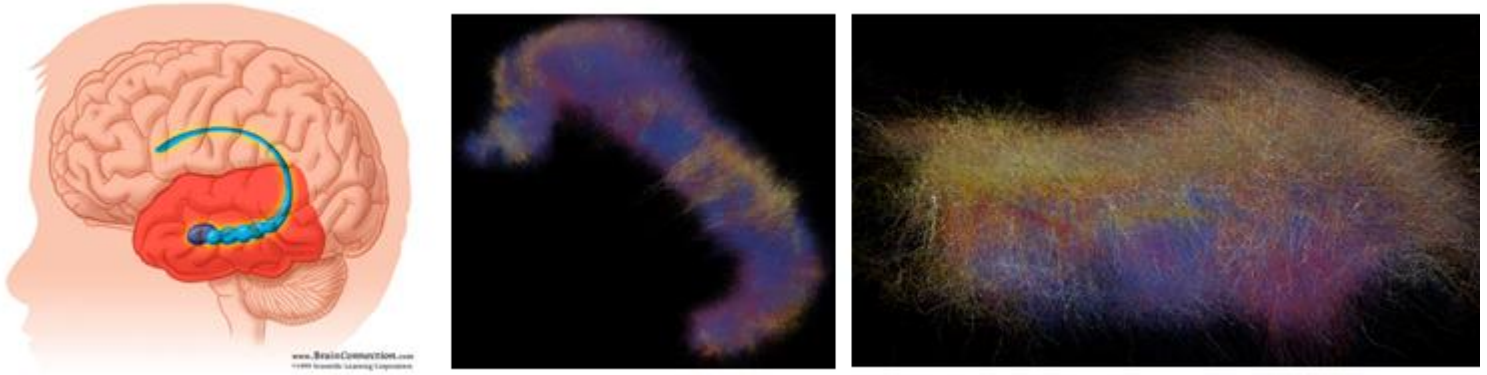

Required resources of the Fenix Infrastructure:

- **Scalable Computing** Services - for running large-scale simulations using Neuron
- **Active Data Repositories** - as temporary storage [write from simulation, read for analysis]
- **Interactive Computing** Services - for analysing data produced by simulations
- **Archival Data Repositories** - for storing final data products

Where to get help

- Central point of contact for any support requests:

  support@humanbrainproject.eu

- Provides access to all levels of support
  - 1st and 2nd level support at the Fenix sites for technical problems with Fenix services
  - Advanced support by the HBP High-Level Support Team (HLST)
- HLST support includes
  - Guidance and assistance in the preparation of applications for computing and data resources
  - Implementation support for simulation and data analytics workflows, possibly using multiple Fenix services, up to the co-development of solutions that benefit other use cases as well

OpenStack SWIFT

- Object Store
  - This is **not** a file system (especially not POSIX)
  - HTTP-based API in the style of REST
  - Manipulation of entities via POST/PUT/GET
  - Atomic units: objects; not files
  - No byte-granularity access
- https://endpoi.nt/my-project/my-account/a-container/object-1/name-1
  - Endpoints signify installations at different sites
  - Projects and accounts group containers by creator
  - Containers group objects into ‘directories’
  - Object names can contain ‘/’, but these hold no meaning
- Sharing Data
  - ACLs can be set for containers
  - Regulate list/read/write access

Archival Data Repository Access

- Currently only at CSCS
  - Obtain an account with the Pollux Infrastructure

    https://user.cscs.ch/tools/openstack/#access-to-pollux
- Install the OS Swift CLI Client

```
# Create a Python Virtual Environment
$ python3 -venv pollux
$ source pollux/bin/activate
# Install the Client
(pollux) $ pip install python-swiftclient python-openstackclient oauthlib python-heatclient
# Get environment script for CSCS Pollux
(pollux) $ curl -O https://raw.githubusercontent.com/eth-cscs/openstack/master/cli/pollux.env
(pollux) $ deactivate
```

- Authenticate

```
$ source pollux/bin/activate
(pollux) $ source pollux/pollux.env
Username: ********
Password: ********
Choose Project: **
# From here you can use the `swift` command line tool.
# `swift` is an alias pre-configured to use your account prefix and the Pollux endpoint.
# ! BE AWARE WHEN TRYING TO ACCESS DIFFERENT ENDPOINTS/ACCOUNTS !
```

FENIX RI

Archival Data Repository Workflow

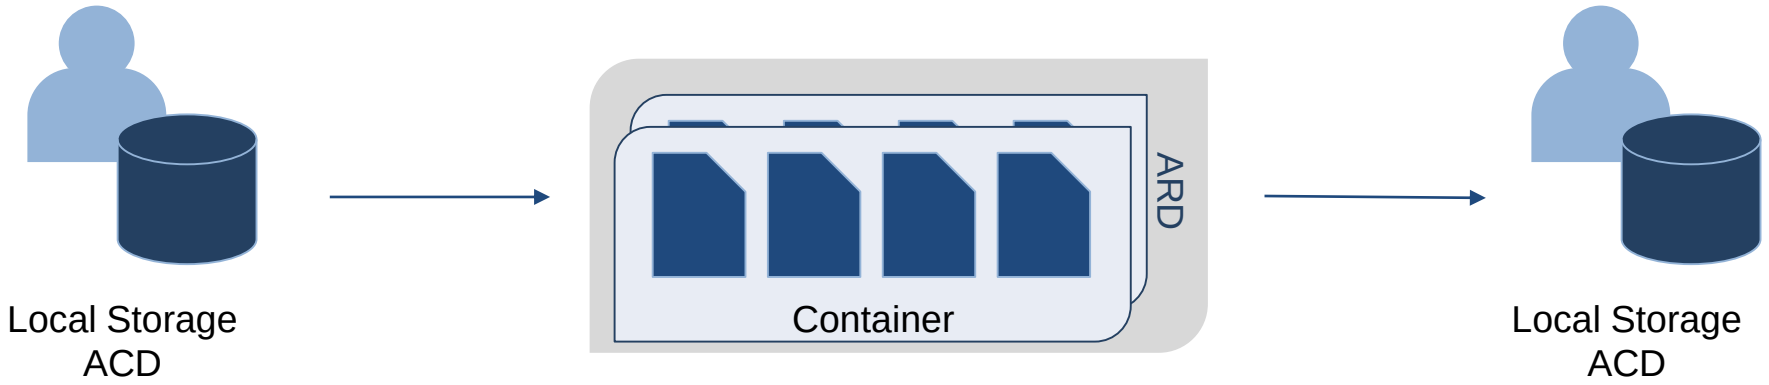

Local Storage
ACD

Local Storage
ACD

```
# Create some data items
(pollux) $ mkdir demo-data
(pollux) $ touch demo-data/a demo-data/b
# Create container `demo`
(pollux) $ swift post demo
# Upload local `demo-data` to `demo`
(pollux) $ swift upload demo ./demo-data
demo-data/b
demo-data/a
# Check container
(pollux) $ swift list demo
demo-data/b
demo-data/a
```

```
# List all containers
(pollux) $ swift list
…
demo
…
# Download the `demo` container
(pollux) $ swift download demo
demo-data/a [...]
demo-data/b [...]
# Add a new object to `demo`
(pollux) $ touch demo-data/c
(pollux) $ swift upload demo demo-data/c
```

There is a lot going on behind the scenes! See here for getting started
https://wiki.humanbrainproject.eu/bin/view/Collabs/how-to-data-access-and-efficient-io/

FENIX RI

ARD - Performance

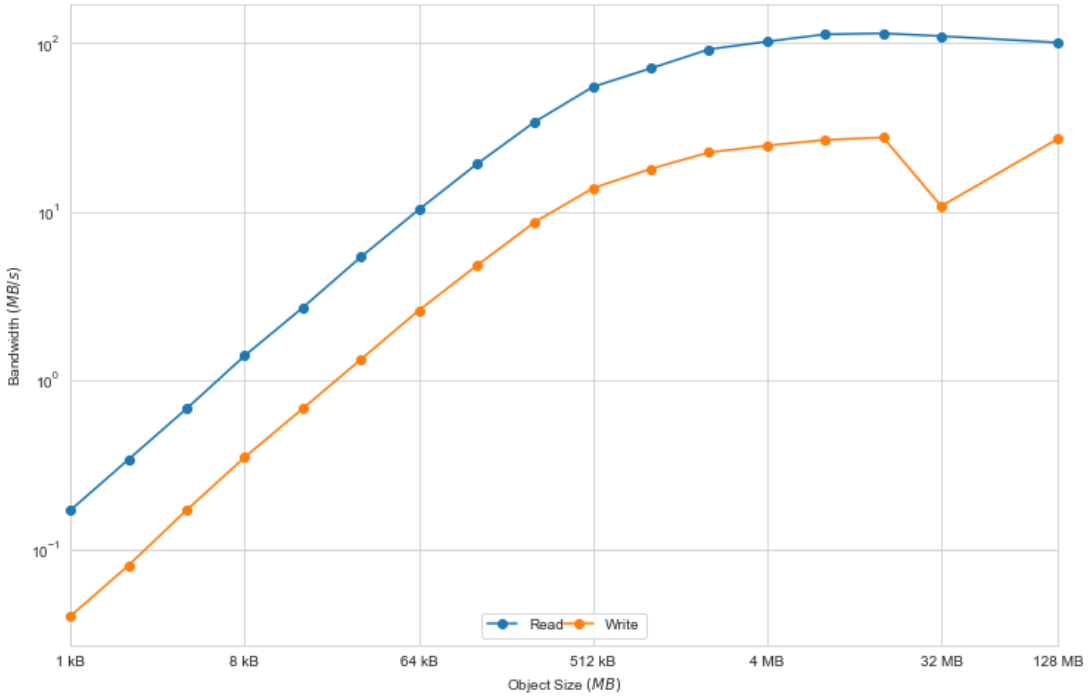

- CosBench (openio 0.4.1) on a single node (https://github.com/open-io/cosbench)
- JSC Juron to CSCS Pollux
- Concurrent mix of 80% read, 20% write
- Object size relevant for performance, use > 4MB

FENIX RI

THANK YOU

fenix-ri.eu
@Fenix_RI_eu
icei-coord@fz-juelich.de