Final release of the federated HPC, Cloud and storage infrastructure for EBRAINS (D6.4 – SGA3)

Figure 1: Federated HPC, Cloud and storage infrastructure

The federated HPC, Cloud and storage infrastructure for EBRAINS is developed, deployed and operated by five European supercomputing centres, which also contribute to the Fenix infrastructure, namely BSC, CEA, CINECA, ETHZ-CSCS and JUELICH-JSC.

Project Number:	945539	Project Title:	HBP SGA3

Document Title:	Final release of the federated HPC, Cloud and storage infrastructure for EBRAINS
Document Filename:	D6.4 (D60) SGA3 M42 SUBMITTED 230831.docx
Deliverable Number:	SGA3 D6.4 (D60)
Deliverable Type:	Report
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA3 M42 / 30 SEP 2023
Actual Delivery Date:	SGA3 M41 / 31 AUG 2023
Author(s):	Anna LÜHRS, JUELICH (P20)
Compiled by:	Anna LÜHRS, JUELICH (P20)
Contributor(s):	Cristiano PADRIN, CINECA (P13), contribution to Sections 2.1.1 and 2.1.2 Hamid FARD, TUDA (P85), contribution to Section 2.1.3 Felix WOLF, TUDA (P85), contribution to Section 2.1.3 Taylan ÖZDEN, TUDA (P85), contribution to Section 2.1.3 Björn HAGEMEIER, JUELICH (P20), contributed to Section 2.1.4 Alex UPTON, ETHZ (P18), contributed to Sections 2.1.4 and 2.2.1 Javier BARTOLOMÉ, BSC (P5), contributed to Section 2.1.5 Benedikt VON ST VIETH, JUELICH (P20), contributed to Section 2.1.6 Kay THUST, JUELICH (P20), contributed to Section 2.1.6 Gilles WIBER, CEA (P11), contributed to Section 2.2.2 Wouter KLIJN, JUELICH (P20), contribution to Section 2.2.3 Marissa DIAZ PIER, JUELICH (P20), contribution to Section 2.2.3 Anna LÜHRS, JUELICH (P20), contribution to Sections 1, 2.2.4, 3, and Annex A Anne NAHM, JUELICH (P20), contribution to Section 3 and Annex A Boris ORTH, JUELICH (P20), contribution to Section 3 and Annex A … and many additional, indirect contributors from the WP6 Tasks developing and operating the services described in this document.
WP QC Review:	Maren FRINGS, JUELICH (P20)
WP Leader / Deputy Leader Sign Off:	Lena ODEN, JUELICH (P20)
T7.4 QC Review:	N/A
Description in GA:	Final release of the improved and enhanced federated HPC, Cloud, storage and network infrastructure and related services for EBRAINS
Abstract:	This document summarises the final release of the improved and enhanced federated HPC, Cloud, storage and network infrastructure and related services for EBRAINS, which are developed in HBP Work Package 6. It particularly focuses on the

	developments pertaining to the central data location and transfer services, a co-scheduling Slurm plugin, federated identity and access management, a monitoring service, and how to enable interactive computing. Moreover, recent activities in the areas of infrastructure operations, policy management and security, of the Scientific Liaison Unit, and infrastructure training and webinars are briefly presented. This document concludes with a brief outlook into planned future activities and challenges.
Keywords:	Infrastructure, Cloud, High-Performance Computing, HPC, storage, network, security, monitoring, co-scheduling, Slurm, data services, interactive computing, Authentication and Authorisation Infrastructure, AAI, infrastructure operations, use case management, user training, Scientific Liaison Unit, FURMS, Fenix
Target Users/Readers:	Developers of platform services, infrastructure users, computer scientists, HPC community

Table of Contents

Table of Figures

1. Introduction

This Deliverable summarises the final release of the improved and enhanced federated HPC, Cloud, storage and network infrastructure and related services for EBRAINS that were developed and operated by the EBRAINS Computing Services Work Package (WP6). An earlier version of this document was submitted as Deliverable D6.2 "Intermediate release of the federated HPC, Cloud and storage infrastructure for EBRAINS"[1].

EBRAINS Computing Services developed, deployed, integrated and operated a variety of basic IT services within the distributed HBP and EBRAINS infrastructure, including in particular the federated HPC, Cloud, storage and network infrastructure, building on resources and services of the Fenix infrastructure that were made available through the ICEI project. More details about the relation of EBRAINS Computing Services to the Fenix infrastructure are included in Annex A: EBRAINS Computing Services and the relation to the Fenix infrastructure which was also already included in the previous Deliverable D6.2, with minor updates for this document. Together with Fenix, EBRAINS Computing Services provided the computing and storage backbone of the EBRAINS infrastructure during the final phase of the HBP.

This part of the project was also in charge of the Neuromorphic Computing systems and services. Please note that the Neuromorphic Computing systems and services are not part of this Deliverable but are described in Deliverable D6.3 "NMC in EBRAINS with TRL8 software". The same applies to the Collaboratory, a workspace in the Cloud for research, development, documentation and collaboration, which is described in Deliverable D6.5 "Final release of the Collaboratory". The Health Data Cloud is covered by the correspondent Deliverable D6.7.

This Deliverable particularly focuses on recent developments and final results with regards to the Central Data Location Service (see Section 2.1.1), the Central Data Transfer Services (see Section 2.1.2), the Co-scheduling Slurm plugin (see Section 2.1.3), AAI and accounting (see Section 2.1.4), the production-level monitoring service (see Section 2.1.5), and making the infrastructure usable for interactive computing (see Section 2.1.6). Moreover, recent activities in the areas of infrastructure operations (see Section 2.2.1), policy management and security (see Section 2.2.2), of the Scientific Liaison Unit (see Section 2.2.3) and related to infrastructure trainings and webinars (see Section 2.2.4) are briefly described. It concludes with an outlook into plans for the future beyond the end of the HBP Flagship (see Section 3).

This document aims at giving an overview of recent developments and activities and refers, where available, to more detailed information about the respective services and developments (e.g. to so-called "Output Documents", websites or repositories). In includes short status updates for all Outputs (OPs), which were delivered in the second half of the current HBP phase.

[1] Available on the HBP website: https://www.humanbrainproject.eu/en/science-development/scientific-achievements/deliverables/third-specific-grant-agreement/

2. Available services and recent activities

2.1 New and updated EBRAINS Computing Services

2.1.1 Central Data Location Service

Short description of the service:

The Central Data Location Service (Outputs OP6.5 and OP6.14) facilitates the end-user's operation to find where data are stored in the ICEI ecosystem and allows the recording of paths (at the same or other sites) for new branches of the same dataset.

Large data sets, that are needed for or generated by EBRAINS workflows, can and should be stored in the Archival Data Repositories. The Data Location service then provides the link between the storage of the data and the access via the data in the Knowledge Graph[2]. By using the Data Location service, the physical location of the data is automatically identified and, in the next step, transported via the Data Transfer service (see Section 2.1.2), that it is now fully integrated with, or using the Data Mover (which has been delivered by the ICEI project to allow the site-local data transfer between object storages and close file systems) to the location where they are processed in workflow. The Data Location service therefore is the interface between the Data Services in Fenix and the Knowledge Graph. This allows the definition of workflows independent of the actual data location.

An overview of the data services available is displayed in Figure 2.

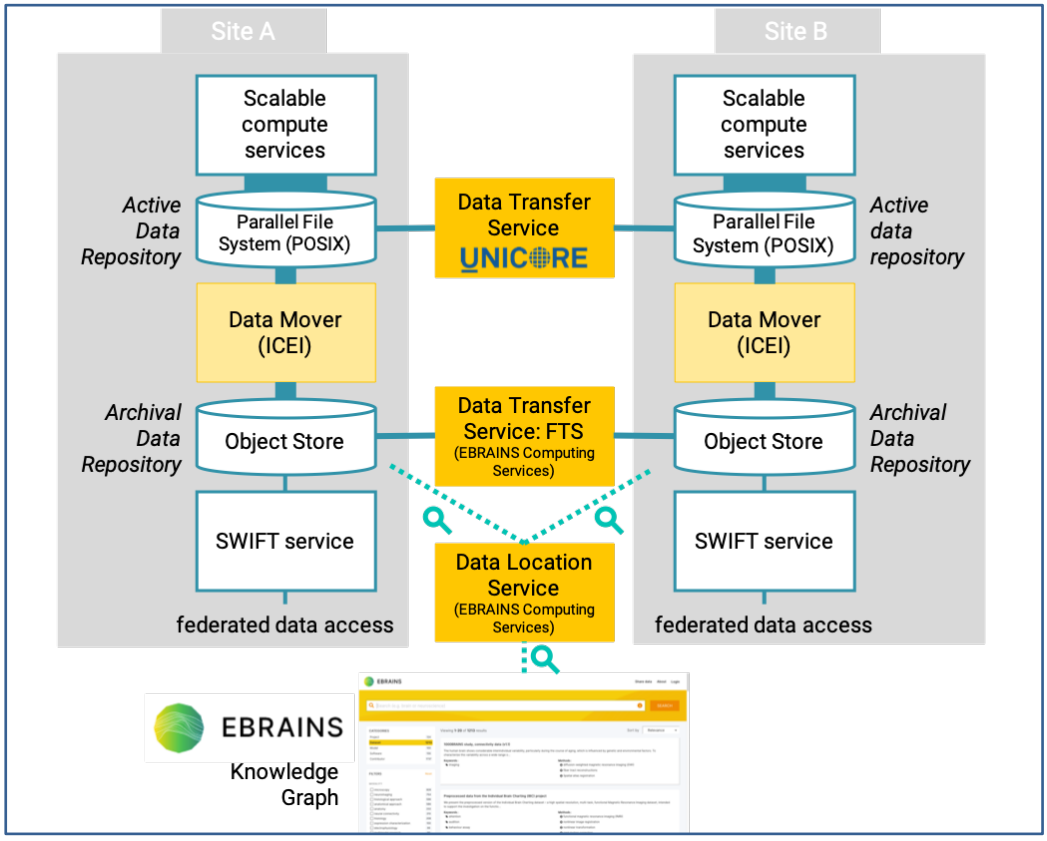

Figure 2: Overview of the Data Services provided by ICEI and EBRAINS Computing Services

Since the Knowledge Graph has a similar functionality for identifying the location of a dataset, it was agreed by Task T6.3 with the Data Curation[3] and Knowledge Graph tasks to focus on the co-

[2] https://search.kg.ebrains.eu
[3] https://www.ebrains.eu/data/share-data/

design of a solution to facilitate the retrieval of data stored in the Fenix infrastructure. The advantage of this solution is that it can also be coupled with the FTS3 data transfer services, in contrast to the solutions readily available in the Knowledge Graph.

The tool used to implement this service is RUCIO, a data management tool developed by CERN, that is fully compatible with FTS3, the technical solution adopted for object-to-object data transfer. The strategy shows that the end-user can update the list of reachable branches in the database of RUCIO and can mirror the produced data set in new branches. The database of RUCIO can be used by the central data transfer service to retrieve the information where data are or where data can be transferred or mirrored.

Documentation:

https://rucio.cern.ch/documentation/

Plans for the future:

Because the HBP community so far has only demonstrated a low interest for this service, and because the openMINDS initiative coupled with the Knowledge Graph can cover the same service, the central data location will probably not be maintained. If there is an interest by other communities to leverage these developments for their infrastructure, e.g. when using the Data Transfer Service, the service could be adjusted accordingly in future projects.

2.1.2 Central Data Transfer Services

Short description of the service:

The Central Data Transfer Services allow end-users to transfer data as object files from/to ICEI object storages (see Figure 2). FTS3[4] is the technology solution adopted for this purpose.

Task T6.3 performed a deep analysis of the use cases provided by the HBP community, to identify the appropriate tool to deploy the central data transfer services. We have chosen FTS3 as a solution, developed by CERN, because it is able to manage data stored both in Swift and S3 object storages, and it is compatible with the most adopted transfer protocols.

FTS3 has been installed and configured at BSC, object storages provided by ETHZ-CSCS, BSC and CEA have been linked in the ICEI federation, and additionally the JUELICH-JSC object storage was linked, even if provided as resource out of ICEI. Task T6.8 took also care to integrate the FTS3 command line interface with the Fenix AAI. From the close collaboration between Task T6.3 and Task T6.8, guides for administrators and end-users have been deployed, like as the deployment strategy documentation. The service has been initially thought to be released with one single instance, but it was then planned to install and configure a second instance in case the number of concurrent usages will exceed the limits set or have an impact on expected performances.

The installation of WebFTS, and its integration with the Fenix AAI was achieved as well, and validation tests have been successfully performed. Finally, as discussed in Marseille during the HBP Summit Week in March 2023, Task T6.3 will explore possible solutions to integrate FTS3 with the Data Mover Service (procured in ICEI), aiming to release a useful plan to identify a workaround waiting for the orchestrator service (developed in "EBRAINS Modelling Services" work package).

A technical summary of the Central Data Transfer Service is available as a confidential report as well (related to Milestone MS6.6[5]).

More information:

https://wiki.ebrains.eu/bin/view/Collabs/ebrains-computing-services-wp6-public/Central Data Transfer Services/

[4] https://gitlab.cern.ch/fts/fts3

[5] This confidential document can be obtained upon request to the EBRAINS Computing Services management: wp6-coord@humanbrainproject.eu

Documentation:

- Documentation of the service, including technology choices made during the development:
  - https://drive.ebrains.eu/f/5257f213815949a89b33/ (Output OP6.8, as of April 2021)
  - https://drive.ebrains.eu/f/853c54a2e2bb4f05a344/ (Milestone MS6.6, as of April 2023)
- User Guide: Swift Transfer with FTS: https://wiki.ebrains.eu/bin/view/Collabs/ebrains-computing-services-wp6-public/Central Data Transfer Services/User Guide%3A Swift Transfer with FTS/ (Output OP6.13)

Plans for the future:

The Central Data Transfer Services could be maintained through future EBRAINS-related projects (like the planned EBRAINS-SERV project[6]), EBRAINS National Nodes or local sites that will provide object storages. The service should be integrated with the Data Mover Service, to allow for a useful workaround waiting for the orchestrator service and support those workflows as service which need data transfer as one of their components.

2.1.3 Co-scheduling Slurm plugin

Short description of the service:

This plugin (Output OP6.10) enables the co-allocation of compute and data resources on a shared multi-tiered storage cluster by estimating waiting times when the high-performance storage (burst buffers) will become available to submitted jobs. Based on the current job queue and the estimated waiting time, the plugin decides whether scheduling the high-performance or lower-performance storage system (parallel file system) benefits the job's turnaround time. The estimation depends on additional information the user provides at submission time.

As follow-up of SGA2 activities, Task T6.3 continued to develop and test the co-scheduling plugin for Slurm on ICEI systems. Because at the beginning of the project no site was able to provide access to the right architecture (multi-tier storage), T6.3 started the development of the initial test environment based on virtual machines using Vagrant and migrated the plugin to MARCONI100. As this system only provides node-local, but the plugin targets remote-shared burst buffer architectures, T6.3 deployed a workaround for the integration based on BeeGFS On Demand (BeeOND), hosted on the NVMe equipped on Marconi100's nodes. Results demonstrated that, based on its estimations, the plugin targeted the correct storage tier, and the integration of the plugin into Slurm on MARCONI100 has been considered successful.

As soon as GALILEO100 has been released, T6.3 started to migrate the test environment on this ICEI HPC system, to replicate the validation phase. This was necessary because GALILEO100 is the only ICEI system equipped with the appropriate hardware to allow the achievement of TRL7 for the plugin, the DDN's IME (Infinite Memory Engine), a well-suited solution for the targeted storage architecture.

More information:

https://github.com/HumanBrainProject/coallocation-slurm-plugin

Documentation:

- https://github.com/HumanBrainProject/coallocation-slurm-plugin/blob/master/README.md
- https://wiki.ebrains.eu/bin/view/Collabs/ebrains-computing-services-wp6-public/Slurm plugin for resource co-allocation/

Plans for the future:

The HBP community has shown a very light interest for this plugin, and we have not found any real HBP use case that can take advantage of the deployment of this service either. While we strongly suggest continuing the investigation about possible real applications (also from other communities),

[6] The proposal for this EU project was submitted in early September 2023.

we do not think that the service should be maintained in the near future. Additionally, the ICEI systems cannot guarantee the right target architecture, and IME, the only one available during the last period of SGA3, has been declared at the end-of-life by DDN.

2.1.4 AAI and accounting services

Short description of the service:

FURMS[7], the Fenix User and Resource Management Service, offers a one-stop-shop for users and administrators of the federated infrastructure, providing management of community and project membership, management of resource allocations for communities and projects, SSH key management, and management of policy documents including paper-based workflows. In addition, it also provides accounting functionality to aggregate accounting information from across the Fenix sites and visualisation of consumed budgets, in addition to budget notification.

FenixAAI, the Fenix Authentication and Authorization Infrastructure, provides users with seamless access to Fenix services and resources. In addition, it supports different levels of assurance (LoA) according to the origin of federated identity providers (IdPs), and is able to grow to incorporate new centres if needed.

While FURMS and FenixAAI are results of the ICEI project, EBRAINS Computing Services have adapted the generic Fenix infrastructure services to the specific needs of EBRAINS (Output OP6.9) and organised a training event for the EBRAINS user community as part of a "Simulate with EBRAINS"[8] training in November 2022 (Output OP6.11).

A short report (related to Milestone MS6.4[9]; Output OP6.9) giving an overview of FURMS and FenixAAI services is available, which specifically shows how a user can access a number of OpenStack systems across the different ICEI/Fenix partners using a single account via the Fenix AAI. It demonstrates how a user can query details about their project via FURMS, such as resource allocation, resource usage and project members.

More information:

FURMS: https://fenix-ri.eu/infrastructure/services/furms

Documentation:

- FURMS documentation: https://unity-idm.github.io/furms/
- Brief instructions for using FenixAAI can be found on the bottom of the following page: https://central-ds.fenix-ri.eu/ds/

Plans for the future:

Discussions around the future of Fenix, including services such as FenixAAI and FURMS, are on-going. In addition to the HBP community, there has been interest in these services from initiatives such as EuroHPC and DestinationEarth, so it is hoped that a way will be found to sustain these services after the end of HBP and ICEI-Fenix. The services is also of high interest for future EBRAINS-related projects, as FenixAAI and FURMS form the basis for a federated infrastructure.

[7] See ICEI Deliverable D3.3: https://fenix-ri.eu/about-fenix/deliverables
[8] https://flagship.kip.uni-heidelberg.de/jss/HBPm?m=showAgenda&meetingID=242
[9] This confidential document can be obtained upon request to the EBRAINS Computing Services management: wp6-coord@humanbrainproject.eu

2.1.5 Production-level monitoring service

Short description of the service:

A production-level service (Output OP6.20) for monitoring the operation of the distributed e-infrastructure services and for incident reporting was put in place based on Icinga[10] and is available[11] to any EBRAINS user. This service is of special interest for platform administrators and users, which directly use infrastructure services. Using this tool, EBRAINS infrastructure users are able to see the status and healthfulness of any of the infrastructure hosts and services of the federated sites. It also keeps a log of any issue detected, which helps to debug any issue or outage, if it was a platform-related or an infrastructure-related cause. The tool permits users to subscribe for notifications via email when a specific infrastructure service is suffering an outage or issue and about its corresponding recovery. Also, an API has been installed so that monitoring information can be consumed by other EBRAINS tools/services.

The service is read-accessible for anyone with an EBRAINS account, while members of the related Task T6.7 have additional privileges on this website. A dashboard (see Figure 3) shows current incidents at service and host level, as well as recently recovered services. The service grid (see Figure 4) gives a good overview which services at which hosting site are at which status.

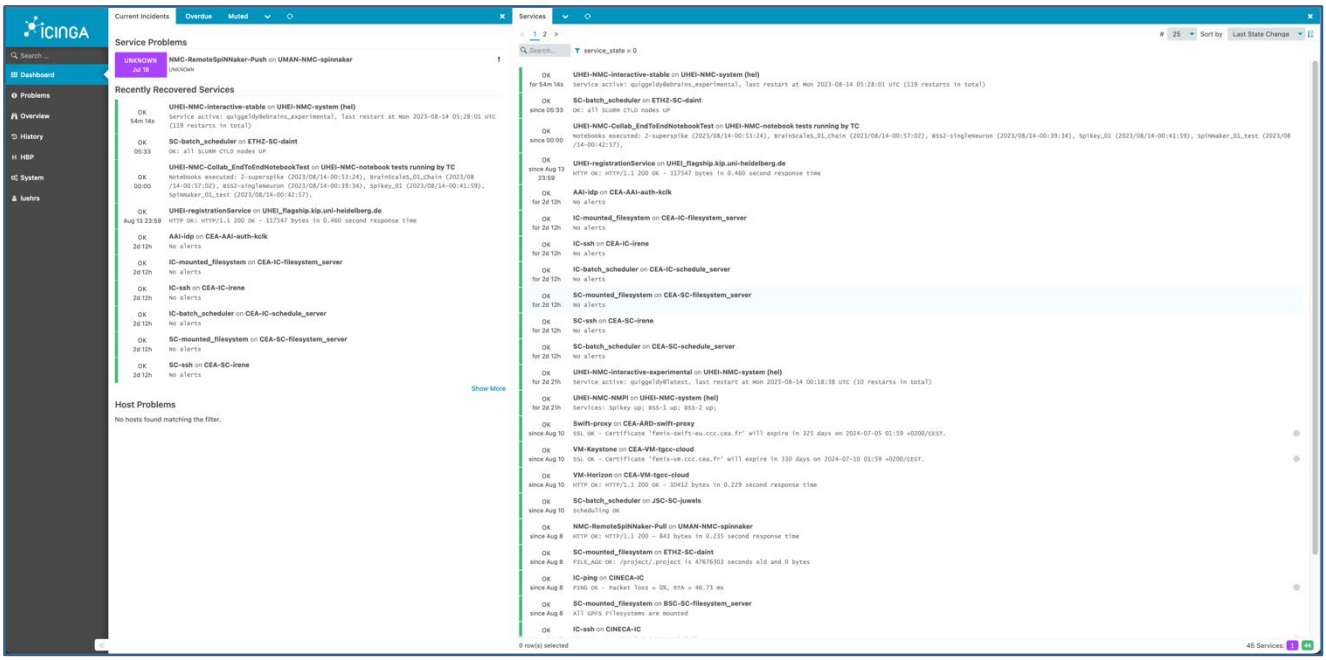

Figure 3: Dashboard of the monitoring service

[10] https://icinga.com/
[11] https://wiki.ebrains.eu/bin/view/Collabs/ebrains-computing-services-wp6-public/Monitoring%20service/

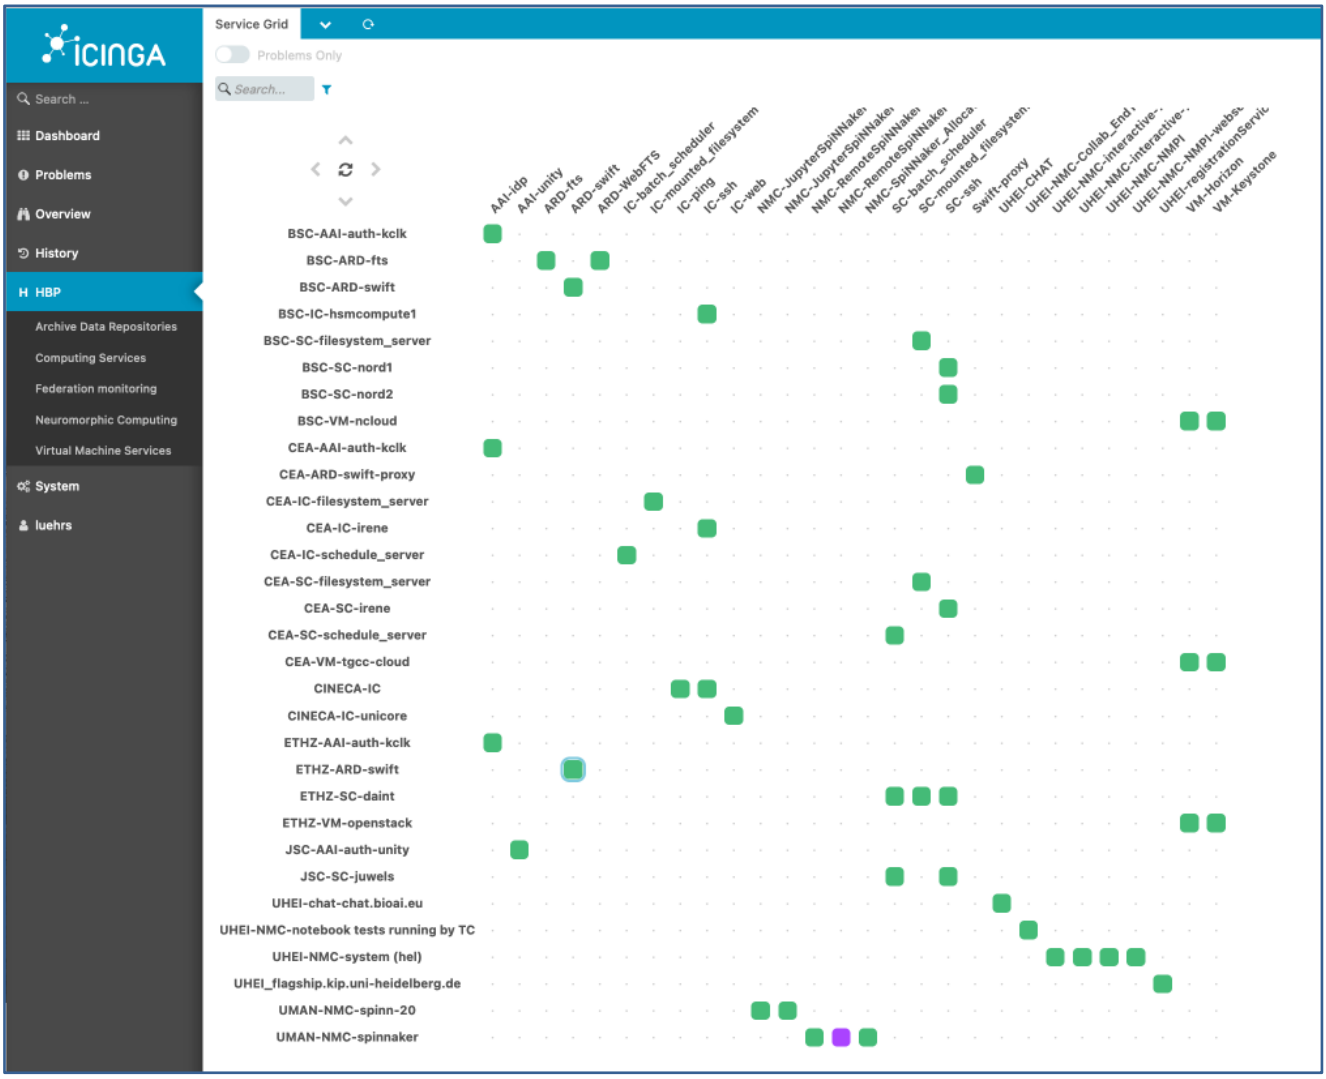

Figure 4: Service grid of the monitoring service

More information:

https://hbpmonitor.bsc.es

Documentation:

- Access to the internal documentation of T6.7 can be provided upon request to wp6-coord@humanbrainproject.eu.
- General documentation about the Icinga tool:
  - https://icinga.com/docs/icinga-2/latest/doc/01-about/
  - https://icinga.com/docs/icinga-web/latest/

Plans for the future:

BSC as leader of Task T6.7 does not participate in the EBRAINS-SERV proposal and is neither part of the Spanish National Node. The monitoring service is composed of a virtual machine, which can be easily transferred and migrated to another of the infrastructure sites in conjunction with credentials and instructions for its operation, so that the service can be maintained by another site in the future.

2.1.6 Infrastructure usable for interactive computing workloads

Short description of the service:

An interactive computing (IAC) service (Output OP6.22) is offered within EBRAINS by the five ICEI sites which also provide the underlying infrastructure as part of Fenix, namely BSC, CEA, CINECA, ETHZ-CSCS and JUELICH-JSC. The intended use cases are workflows which have high resource demands but also require quick feedback loops. Typical examples are pre- and post-processing of data for simulations on the scalable computing services (SCC) and generally analysis of data volumes which is not feasible on the user's local machine.

All systems deployed as part of the Fenix infrastructure provide generic hardware and software capabilities that can handle a sufficiently wide range of usage scenarios well. All systems offer GPUs for visualisation and computing purposes. The IAC resources at BSC and CEA are particularly well suited for lowly-scalable applications requiring a large node-addressable memory space. Most IAC services are collocated with resources designed for batch-processing, i.e. SCC. The software used at the sites to interact with the deployed IAC resources is comparable, but there are relevant differences in particular with regards to the scientific software stacks.

The IAC service is designed based on an in-depth analysis of use cases gathered and described in the public ICEI Deliverable D3.6 "Scientific Use Case Requirements Documentation"[12]. In contrast to previous assessments, this analysis was focused on the extraction of information for the configuration and optimisation of the deployed systems rather than the development of future components, because the EBRAINS Computing Services Task in charge (T6.9) of the IAC service is mainly focused on operations.

While from a user's perspective there is a large overlap of IAC and SCC, the IAC service can be tuned separately to accommodate specific requirements of the use cases. The IAC resources at all sites use the workload manager Slurm for job scheduling and resource management. The heterogeneous job feature of Slurm supports the co-scheduling of SCC and IAC jobs, so that both resources are available at the same time. Spawned processes of the job are also placed in the same process management namespace, which allows MPI applications to use intra- and inter-communicators to send messages between both types of resources.

An important goal of the interactive computing service is to speed up the resource allocation to guarantee interactivity for the user in contrast to a batch-based usage mode. The approach taken to achieve this is to analyse the batch system behaviour and to propose operation enhancements of the batch system Slurm. Aspects to work on in particular are reservations, separate job partitions as well as job boosting and pre-emption.

We have considered a number of different options to enable the IAC service at all sites, including

- Workload manager (Slurm) configurations: scheduler parameters, resource dedication (at some sites IAC and SCC share the same resources, which can lead to competing job allocations);
- Priority shaping: prioritising interactive jobs over other pending requests, in particular SCC jobs;
- Alternative operational models: pre-emption (premature halting or termination of running jobs to free up resources for recently submitted interactive job requests), resource oversubscription.

While the workload manager is identical at all sites, there are many either technical or policy level boundary conditions that have an impact on the way the IAC service can be offered across them.

While, for example, job boosting is an easy means to speed up the scheduling of jobs, it makes the scheduling for other users of the systems unpredictable and is therefore not supported everywhere. Another example is pre-emption, which requires adjusting the jobs that are cancelled in favour of interactive jobs, because without taking snapshots a cancelled job could potentially lose a significant amount of computed results. Due to these boundary conditions and the varying offers of the IAC

[12] https://fenix-ri.eu/about-fenix/deliverables

service, the use cases requesting it need to be well understood to offer the best possible implementation.

Although different approaches for reliably short waiting times were considered many proved to still exceed the threshold of a few minutes. The three contradicting main requirements are: system utilization, waiting time (for jobs to start) and reliability (of performance and jobs not being stopped). Since IAC usually involves a scientist actively waiting for a response of the system, it is very sensitive to waiting times. Therefore, the system utilization was prioritized lowest of the three requirements, meaning there are dedicated compute resources potentially left waiting for the next interactive requirement so they can be immediately be processed.

For workloads which need few compute resources but might rely on data residing in one of the HPC sites, all sites provide login nodes which usually provide more compute capabilities and memory than the user's notebook. They are also close to the data with high bandwidth and low latency.

More demanding use cases are workloads requiring multiple nodes. Those are handled by dedicated systems (BSC) or partitions (CEA, CINECA, CSCS, JSC). Most partitions are dedicated and designed for only a few nodes. If larger jobs are required, the SCC generally support interactive use, although it usually does not provide the short waiting times.

The IAC service also enables access to compute resources orchestrated from a web service, which itself is intended to provide an interactive experience to its users. The integration of a web service or other platforms is possible using UNICORE, based on work in earlier phases of the HBP. The IAC workload manager and a UNICORE TSI daemon are integrated via the UNICORE RESTful API, which makes CPU and GPU resources accessible from a web service. Reverse connections using https or secure shell protocols to select platform hosting servers are possible from the login partition of the IAC systems.

Visualisations are an important aspect pointed out in the use case descriptions. Since community-specific software is usually not pre-installed and can even be relatively difficult to install, it is possible to ship customised software environments in the form of container images and to execute them within the container runtime environments offered at all sites. Most sites also offer JupyterHub services, including browser-based visualisation, for easy access and a flexible and sophisticated interface to IAC and SCC.

More information:

https://fenix-ri.eu/infrastructure/resources/available-resources

Documentation:

- BSC: https://www.bsc.es/supportkc/docs/HSMCompute/intro
- CEA: https://www-hpc.cea.fr/tgcc-public/en/html/toc/fulldoc/Interactive_access.html#remote-desktop-system-service-nicedcv
- JSC: https://docs.jupyter-jsc.fz-juelich.de/github/FZJ-JSC/jupyter-jsc-notebooks/blob/documentation/index.ipynb

Plans for the future:

Within the HBP the details of what are the relevant aspects of interactive computing were analysed. Where possible, the services will stay available on a best effort basis. Furthermore, the insights gained are included the design of future of future systems, to cover a broader range of workflows from different communities. Tools like UNICORE and Containers are important for future workflows and therefore should be further maintained.

2.2 Recent and ongoing activities

2.2.1 Operation of HPC, cloud, storage and network services

Short description of the activity:

This activity concerns the operation, maintenance and support of systems providing network, scalable compute, cloud services and archival data repositories provided by ICEI by the five sites across Europe (Outputs OP6.12 and OP6.28). In order to ensure close collaboration with the users of the infrastructure, all sites participated in the weekly HBP Technical Coordination meetings, as well as offering extended support for key HBP use cases and EBRAINS platforms such as Cellular Level Simulation [13], Neurorobotics Platform (NRP) [14], Medical Informatics Platform (MIP) [15], Human Intracerebal EEG Platform (HIP)[16], The Virtual Brain (TVB)[17], ilastik[18], KnowledgeGraph[2] and the Collaboratory[19].

The operation of the interactive computing services is taken care of by Task T6.9 (see Section 2.1.6). Already in earlier phases of the HBP, a ticketing system was put in place and operated by BSC for handling support requests related to the High Performance Analytics and Computing (HPAC) Platform, which is now part of EBRAINS Computing Services. This ticketing system is still in place and connected to the EBRAINS ticketing system, so that all tickets concerning the HPC, cloud, storage and network services are forwarded to the HPAC ticketing system and handled there. Information on the ticket status are sent back to the EBRAINS ticketing system.

More details about the operation of HPC, cloud, storage and network services can be found in a public version of the report "Operational teams defined and operational processes defined" [20] (related to Milestone MS6.1).

Another important activity was the deployment of the EBRAINS software stack[21] across sites. The software stack was defined by the Technical Coordination team at ATHENA, and those parts relevant for HPC environments were then installed at all HPC sites and made available to the user community as modules. Documentation is available in a public collab (see below). More information is also provided in the public Deliverables D6.3 "NMC in EBRAINS with TRL 8 software" and D5.7 "EBRAINS Infrastructure".

More information:

https://fenix-ri.eu/infrastructure

Documentation:

- BSC user guide: https://www.bsc.es/user-support/nord3.php
- CEA user guide: https://www-hpc.cea.fr/tgcc-public/en/html/tgcc-public.html
- CINECA user guide: https://wiki.u-gov.it/confluence/display/SCAIUS/UG3.3%3A+GALILEO100+UserGuide
- CSCS user guide: https://user.cscs.ch/
- JSC user guide: https://apps.fz-juelich.de/jsc/hps/jusuf/index.html

[13] https://www.ebrains.eu/modelling-simulation-and-computing/simulation/cellular-simulation
[14] https://www.ebrains.eu/modelling-simulation-and-computing/simulation/neurorobotics-platform
[15] https://www.ebrains.eu/tools/medical-informatics-platform
[16] https://www.ebrains.eu/health-research-platforms/health-platforms/work-with-health-data-2
[17] https://www.ebrains.eu/tools/the-virtual-brain
[18] https://www.ebrains.eu/tools/ilastik
[19] https://wiki.ebrains.eu
[20] https://drive.ebrains.eu/f/9e3f20bd4889476a9309/?dl=1
[21] https://gitlab.ebrains.eu/technical-coordination/project-internal/devops/platform/ebrains-spack-builds (EBRAINS account required for access)

- Documentation of EBRAINS Software Stack on HPC: https://wiki.ebrains.eu/bin/view/Collabs/ebrains-software-stack/Usage%20on%20EBRAINS%20HPC%20sites/

Plans for the future:

Discussions around the future of Fenix are on-going, as well as the provision of infrastructure to EBRAINS post-HBP. CEA, CINECA and JSC have joined the EBRAINS-SERV proposal (submitted in September 2023), and would make cloud and storage resources available to the EBRAINS community through this project. A possible involvement of CSCS was still under investigation at the time of writing this section.

Recently efforts were started to deploy all services hosted at CSCS also at other sites, starting with the JSC, to be prepared for and more flexible after the end of the SGA3 and ICEI projects.

2.2.2 Policy management and security

Short description of the activity:

An important activity was the definition of a trust zone between the involved infrastructure partners, namely all ICEI sites (BSC, CEA, CINECA, ETHZ-CSCS and JUELICH-JSC) and the neuromorphic sites UHEI (hosting the BrainScaleS system) and UMAN (hosting SpiNNaker). CSC, which is part of the Fenix infrastructure, was included as a second step as well. The trusted zone way is a direct follow-up of the security actions (based on HBP security policies reviewed) done at the beginning of 2020, i.e. in the previous project phase (SGA2). In order to define this trusted zone, and because a simple external security audit cannot be done on each site, all sites decided to complete security self-assessments related to the SCI-V2 document of the WISE community[22] and the Fenix Security Measures Catalogue. The results of the latest iteration of the self-assessments, more details about the established Security Incidence Response Team (C-SIRT) and contributions to related activities in collaboration with other groups within the Human Brain Project are available in a confidential report, which is an update of the confidential report of Output OP6.1.

Plans for the future:

As security should be a key aspect of any federated infrastructure, the self-assessments should be continued in the EBRAINS-SERV project.

2.2.3 Scientific Liaison Unit

Short description of the activities:

In the previous reporting period, the Science Liaison Unit's (SLU) primary focus was on gathering requirements from and working together with the showcases in WP1-3 which were taken up by WP4-6 towards the implementation in the infrastructure (Outputs OP6.32 and OP6.33).

The SLU continued its work to systematically bring the EBRAINS services together. Trough the creation of material including videos, posters, documented processed and other coordination actions, the SLU mapped the infrastructure and services to share it with the researchers, generated descriptions and liaise with researchers to find the best way to integrate their research on EBRAINS. The SLU designed and distributed tools to help researchers find their requirements, and guided the PI's and their groups towards a homogenized delivering of the showcases. A lot of information, activities and results are summarised in the SLU's Collab[23].

While still doing the basic coordination work and processing of requests for information, the SLU put a bigger focus on the needs of the broader science community in which the HBP is embedded. The following major efforts were undertaken:

[22] https://wise-community.org/
[23] https://wiki.ebrains.eu/bin/view/Collabs/ebrains-scientific-liaison-unit-slu/

1) The SLU co-wrote the 3rd booklet on the achievements of the HBP
2) The SLU organized in collaboration with the HBP Education Programme a workshop focused on EBRAINS science workflows and produced a white paper based on the results.
3) The SLU performed a number of surveys of the neuroscience community regarding the usage of the tools.
4) Finally, the SLU provided high-level aggregate information and content for the strategic planning of the future of the EBRAINS RI beyond the end of the HBP.

3rd Booklet on the achievements of the HBP

The Task Force for Science Communication and the SLU wrote the 3rd HBP booklet[24] on the achievements of the HBP: "An extensive guide to the tools developed". In addition to defining scope and content, the SLU collected details on a total of 158 tools developed in the HBP, and, after identification of tools owners, compiled a short description and categorization for each tool in an iterative process. Complemented with success stories from the showcases and testimonials from researchers at all stages of their science careers, this booklet spans 110 full colour A4 pages and was presented to the community in September 2023.

Workshop on workflows

Acknowledging the importance of early career researchers for the continued legacy of the HBP, the SLU organized a workshop on developing workflows with EBRAINS tooling in collaboration with the HBP Education Programme. During the 2-day online event, the students formalised their own scientific ideas using a structured process and mapped these to the tools available in EBRAINS. In the subsequent months a white paper was drafted in collaboration with these students and finally published in the Collab of the SLU: https://wiki.ebrains.eu/bin/view/Collabs/ebrains-scientific-liaison-unit-slu/

Surveys

As one of the potential methods to address the needs of the wider community, the SLU created a number of online surveys to gather input from the HBP community. Two surveys implemented as webforms were sent to the community at large. Although it is impossible to completely prevent any bias in the sample, the aim was to be as inclusive as possible. Requests for participation were sent to the EBRAINS consortium, as an extra HBP newsletter, as part of a normal HBP newsletter, and additionally the National Nodes and EBRAINS Science and Technology Committee (ESTC) representatives were invited to forward the surveys to any potential mailing lists or communication channels targeting scientists that might be interested in having their voice heard.

Standard information points like nationality, affiliation and role complemented the usage information gathered with the surveys. The collected data was sanitised by removing data points for entries where participants did supply any information, and personal information was removed before presenting the survey results as an interactive web applet enabling exploration and interpretation of the results. The surveys saw a total of 430 and 110 participants, respectively, from 34 countries including 8 from outside of Europa.

The results of these surveys are an important contribution for the strategic planning of the future of the EBRAINS RI beyond the end of the HBP projects, as they were already used for the preparation of the EBRAINS-SERV proposal.

More information about the surveys as well as statistics and results can be found in a public Collab of the SLU: https://wiki.ebrains.eu/bin/view/Collabs/slu-hbpb/

International Partnerships and Industry Collaboration

In pursuit of sustainability, the SLU expanded its efforts beyond the EU region, seeking partnerships and collaborations with industries and external stakeholders. With the collaboration of Work Package

[24] http://humanbrainproject.eu/tools

8 and the HBP Partnering Projects, the SLU contributed to the edition of the compilation publication on eNeuro[25].

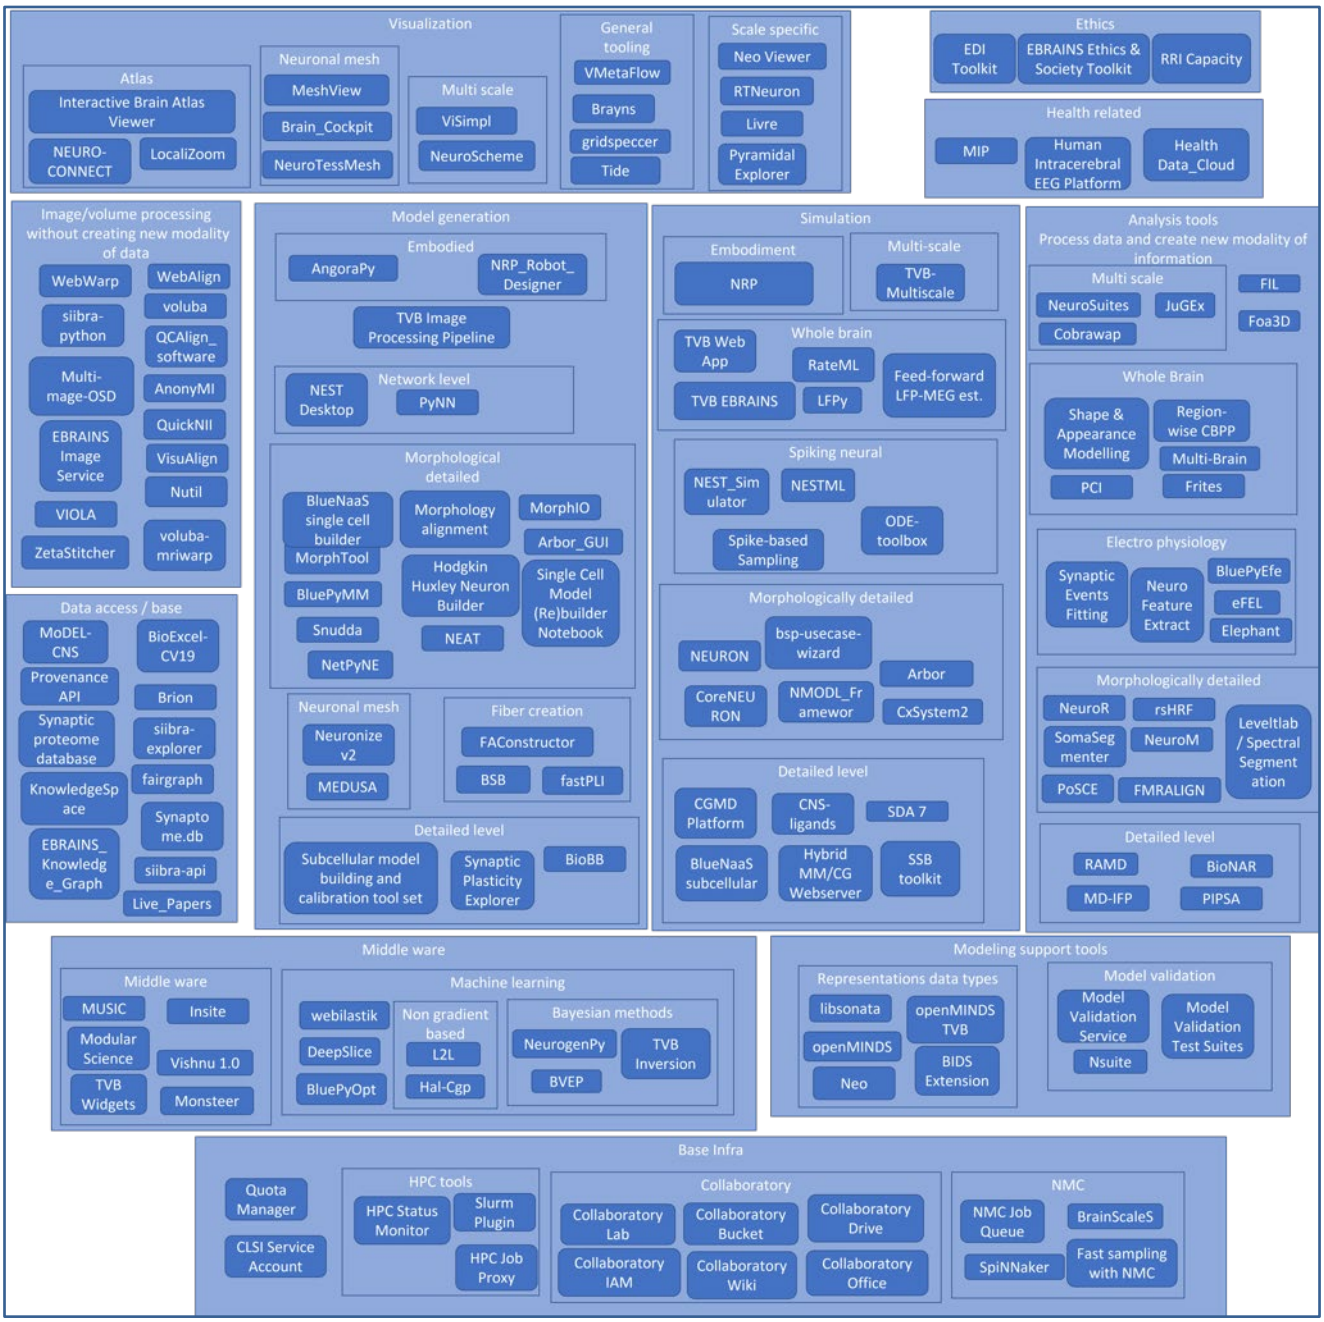

Figure 5: Tools of the HBP put in context with each other

High-level aggregate information and support towards the future of the EBRAINS RI

Due to its role in information gathering regarding the usage of tools and services in HBP and EBRAINS, the SLU has a unique helicopter view of the different activities performed in this project. To support the strategic planning for the EBRAINS RI after the end of the HBP, the SLU performed a summarization and aggregation of information sources and supplied this, e.g., to the writing teams of the EBRAINS-SERV proposal. Examples include the creation of different diagrams displaying the clustering of the tools and activities in the HBP, leveraging machine learning and manually approaches. Additionally, the SLU leveraged its experience in workflow formalisation and co-designed a set of science workflows showcasing the application of EBRAINS tools to new scientific challenges. An example figure as produced in these efforts can be seen below (see Figure 5), displaying the 158 tools of the HBP in context with each other.

[25] Not yet published at the time of writing this section.

Plans for the future:

The SLU plans its continuation in the post-HBP landscape as the Science Support Team (SST) part of the broader support activities in EBRAINS. The SST complements reactive support activities with forward looking activities, intercepting, analysing and supporting users that want to apply EBRAINS to perform novel scientific endeavours. The SST is one of the formal instruments for capturing and conveying new science requirements for EBRAINS, to ensure its continuous support for excellent science.

2.2.4 Infrastructure trainings and webinars

Short description of the activity:

EBRAINS Computing Services, often in close collaboration with the ICEI project, organised several infrastructure trainings, workshops, webinars and other dissemination activities to promote the infrastructure services available. To give some examples:

- Workshop: “Let's explore EBRAINS! - And discover how your scientific use case can be realised” during the 6th HBP Student Conference, 23 February 2022, online
- Talk: “Data Transfer Services in Fenix” during the 16th Tea & Slides event, 3 March 2022, online
- Parallel session about “Massive Computing applied to EBRAINS” at the HBP Partnering Projects Meeting, 5-7 September 2022, Nijmegen, Netherlands
- Infrastructure training event: “Simulate with EBRAINS - How to use EBRAINS research infrastructure and Fenix resources for simulations on a scale from molecules via small and large networks of point or structured spiking neurons to simplified whole brain activity and virtual environments”, 7-10 November 2022, online, with various relevant sessions related to Fenix AAI and FURMS as part of an integrated workshop
- Workshop: “Tools and workflows for realizing your scientific use cases in EBRAINS” during the 7th HBP Student Conference, 19-20 January 2023, Madrid, Spain
- Booths and sessions at the HBP Summit 2023, 28-31 March 2023, Marseille, France

Also, as part of the Fenix Webinar series[26], some of the editions were prepared as a collaboration between EBRAINS Computing Services and the ICEI project.

The SLU actively participated in conferences, workshops, and events to promote EBRAINS within the scientific community and among young researchers. This outreach strategy contributed to increased awareness of the platform's capabilities and its potential to drive cutting-edge research, also expanding its efforts beyond the EU region, seeking partnerships and collaborations with industries and external stakeholders in collaboration with other HBP groups. The SLU contributed to the edition of the compilation publication on eNeuro[25].

[26] https://fenix-ri.eu/media/webinars

3. Looking Forward

In the past years, the growing HBP and EBRAINS community was in the comfortable situation to have access to a dedicated, federated HPC, Cloud and storage infrastructure based on resources provided by the Fenix RI through the ICEI project. The HPC, cloud and storage systems as well as some of the services were procured by the ICEI project, and the EBRAINS Computing Services operated these services for the EBRAINS community and took care of necessary adaptations and developments to customise these generic services for the neuroscientific requirements of the EBRAINS RI. A community call for resource access was continuously open and projects could request any combination of resources of different systems at one or multiple Fenix sites with a single, comparably short proposal.

The situation after the end of the HBP and the ICEI project, which both end in September 2023, will be more challenging once the systems procured by the ICEI project have reached their respective ends of life. Most partners of the EBRAINS Computing Services (and of the ICEI/Fenix consortium) have a strong interest to continue their efforts by further maintaining the established services for the EBRAINS RI, e.g., through contributions to the respective EBRAINS National Nodes. In cases where partners decide to discontinue their EBRAINS involvements in the future, other partners of the federation are preparing to take over the services previously operated by these partners if the services are still required by the EBRAINS RI.

While the ICEI project allowed to procure resources and services at multiple sites at once, the landscape will be a lot more diverse in the future. Now and in the foreseeable future, there is no call on the horizon that could fund the next generation of Fenix systems at once. Therefore, they will need to be acquired through multiple sources for the EBRAINS RI, such as national funding (e.g., as contributions to the respective EBRAINS National Nodes) and maybe partially also through future European funding. Future project proposals could also include funding for infrastructure resources, e.g., to extend existing systems with a then guaranteed partition for the EBRAINS RI. However, this will likely result in a situation where the lifetimes of resources available to EBRAINS users will differ significantly, so that the EBRAINS RI and its users have to plan with and be prepared for systems offered at different sites that have to be replaced at different points in time. Also, some resources might only be available to a subset of the user community if, e.g., nationally funded systems with specific eligibility criteria are integrated.

In the preparation of the EBRAINS-SERV proposal it is planned that CEA, CINECA and JSC provide cloud and HPC resources (and storage resources that come with HPC/Cloud ones) in-kind to the EBRAINS RI during the lifetime of this project. Additionally, they aim for submitting a proposal to request special access to resources on EuroHPC systems to continue the legacy of the HBP, as the HBP is listed among the “initiatives considered as strategic for the Union”[27] by EuroHPC. All Fenix partners located in the EU host a EuroHPC-funded pre-Exascale system[28] or, in case of CEA[29] and JSC[30], will host the first European Exascale supercomputers.

While the path to securing HPC resources for the EBRAINS RI seems to be clear, the work in HBP and in particular the preparation of the EBRAINS-SERV proposal have shown that cloud resources, which are strongly linked to the HPC systems, are important for EBRAINS. While EuroHPC offers the opportunity to apply for the required HPC resources, similar concepts for Cloud resources do not yet exist. Other European Cloud infrastructures at other locations are not a solution for many use cases of HBP and the EBRAINS RI, where, e.g., a service running in the Cloud needs to access the high capacity storage systems connected to an HPC system (e.g., to view atlas data).

It will be important, yet challenging to develop a roadmap for the base infrastructure of EBRAINS that encompasses the diversity with regards to the availability of future infrastructure resources,

[27] https://eurohpc-ju.europa.eu/system/files/2022-06/Decision%2002.2022%20-%20Special%20Access_0.pdf
[28] https://eurohpc-ju.europa.eu/supercomputers/our-supercomputers_en
[29] https://eurohpc-ju.europa.eu/jules-verne-consortium-will-host-new-eurohpc-exascale-supercomputer-france-2023-06-20_en
[30] https://eurohpc-ju.europa.eu/one-step-closer-exascale-eurohpc-ju-and-forschungszentrum-julich-sign-hosting-agreement-exascale-2022-12-14_en

which will naturally also lead to a higher level of diversity with regards to technical specifications. This is a planned key activity in the EBRAINS-SERV project.

Annex A: EBRAINS Computing Services and the relation to the Fenix infrastructure

To make this document more comprehensible and self-contained, this section, with minor updates, is taken from Deliverable D6.2 "Intermediate release of the federated HPC, Cloud and storage infrastructure for EBRAINS" and added again for reference. Updates are indicated below. Tenses were adjusted as needed as well.

The EBRAINS Computing Services Work Package (WP6) developed, deployed, integrated and operated a variety of basic IT services within the distributed HBP and EBRAINS infrastructure, including

- the HPC/Cloud computing and storage services of the Fenix infrastructure[31], which are made available through the ICEI project;
- the neuromorphic computing services SpiNNaker and BrainScaleS;
- the HBP Collaboratory;
- the EBRAINS accounts system ("identity provider", i.e. single sign on for EBRAINS services); and
- the Health Data Cloud (*added to WP6 after submission of D6.2 as a result of an open call*).

This document focuses on the first point, i.e. advances of the federated HPC, Cloud, storage and network infrastructure services. Developments of the neuromorphic systems and related services are described in Deliverable D6.3 "NMC in EBRAINS with TRL8 software", the Collaboratory and EBRAINS accounts system are covered in D6.5 "Final release of the Collaboratory" and the Health Data Cloud is presented in the correspondent Deliverable D6.7.

The systems and services operated by WP6 enable the platform services layer and individual vertical solutions to integrate different EBRAINS services within complex workflows. The services of this infrastructure layer (see Section 2) thus served within the HBP and EBRAINS as a basis for the EBRAINS Data Services[32], EBRAINS Brain Atlas Services[33] and EBRAINS Modelling and Simulation Services[34] and could, of course, also be used directly by internal and external end users. The work was mainly focused on the operation of the infrastructure services to ensure a high quality of service that could be achieved, e.g. by providing robust operational environments and establishing mechanisms that allow for timely identification of problems.

Another important focus was the adaptation of the generic Fenix infrastructure services to the specific needs of EBRAINS, e.g. regarding Authentication and Authorization Infrastructure (AAI) integration and the Fenix User and Resource Management Service (FURMS). Within the Fenix infrastructure, six European supercomputing centres, namely BSC[35] (Spain), CEA[36] (France), CINECA[37] (Italy), CSC[38] (Finland), ETHZ-CSCS[39] (Switzerland) and JUELICH-JSC[40] (Germany), have agreed to align their services. The distinguishing characteristic of this e-infrastructure is that data repositories and scalable and interactive supercomputing systems are in close proximity and well-integrated. An initial version of this infrastructure was realised by BSC, CEA, CINECA, ETHZ-CSCS and JUELICH-JSC through the ICEI project (Interactive Computing E-Infrastructure), which is part of the European Human Brain Project[41] (HBP) and ending at the same time as the HBP SGA3 project. Figure 6 displays

[31] https://fenix-ri.eu/
[32] https://www.ebrains.eu/data/find-data (*updated link*)
[33] https://www.ebrains.eu/brain-atlases/reference-atlases (*additional reference*)
[34] https://www.ebrains.eu/modelling-simulation-and-computing/simulation (*updated link*)
[35] https://www.bsc.es/
[36] http://www-hpc.cea.fr/index-en.htm
[37] https://www.cineca.it/en
[38] https://www.csc.fi/
[39] http://www.cscs.ch/
[40] https://www.fz-juelich.de/ias/jsc
[41] https://www.humanbrainproject.eu/

the relation between site-local and federated infrastructure services, EBRAINS platform services like data and modelling services, and the user communities.

While the ICEI project focused on the procurement of hardware and developments of generic (i.e. not community-specific) solutions such as FURMS, EBRAINS Computing Services operated these services and linked them to the EBRAINS infrastructure by developing, enhancing and operating additional services as required by the other parts of the HBP. This document describes the recent advances in the development areas and additional important activities.

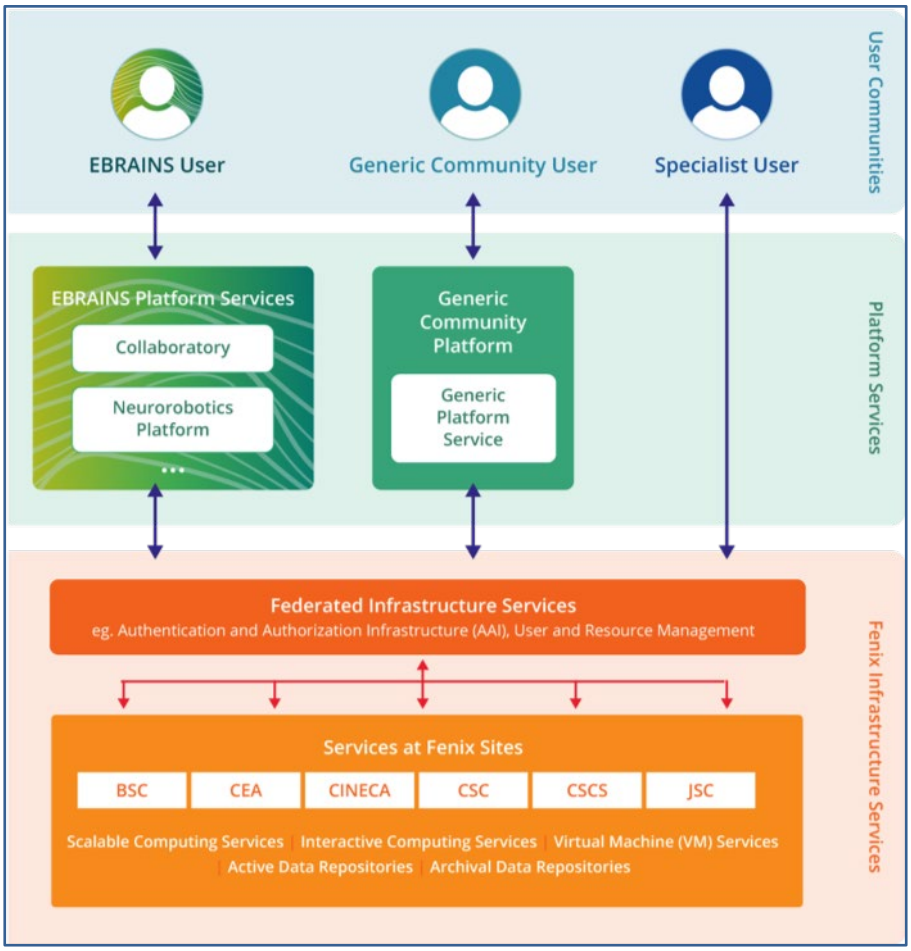

Figure 6: Relation between the Fenix infrastructure services, platform services and user communities

The Fenix sites offer various types of local infrastructure services, on top of which federated infrastructure services are deployed, which are operated by EBRAINS Computing Services. EBRAINS Platform Services run on the infrastructure services and are used by EBRAINS users. This concept can be generalised for other communities as well. Specialist users can also directly use the federated and site-local infrastructure services.

While the design of the Fenix infrastructure was to a large extent based on use cases of the HBP, collected and analysed in the ICEI project, the site-local and federated infrastructure services (i.e. all services displayed in orange) were designed to be community independent, so that they could serve multiple communities. These communities, with HBP and EBRAINS being the first large user community, then develop community-specific platform-level services, which are operated on top of the generic infrastructure services. EBRAINS Computing Services supported the development of interfaces between the generic infrastructure layer and other EBRAINS platform services with the services and activities described in this document.

To date, EBRAINS represents the only large user community, which has developed and deployed their own platform services, but several projects from different communities also successfully applied for ICEI resource allocations via the PRACE-ICEI calls[42]. Moreover, the ICEI project has established contacts with potentially interested communities. The terms "Generic Community User" and "Generic Community Platform" indicate that the Fenix infrastructure is by design not limited to the specific requirements of EBRAINS but can also serve other research communities.

[42] https://fenix-ri.eu/access (see "European Calls for Scientists of all Research Domains")