SP7 High-Performance Analytics and Computing Platform – Results for SGA2 Year 2
D7.6.2 – SGA2

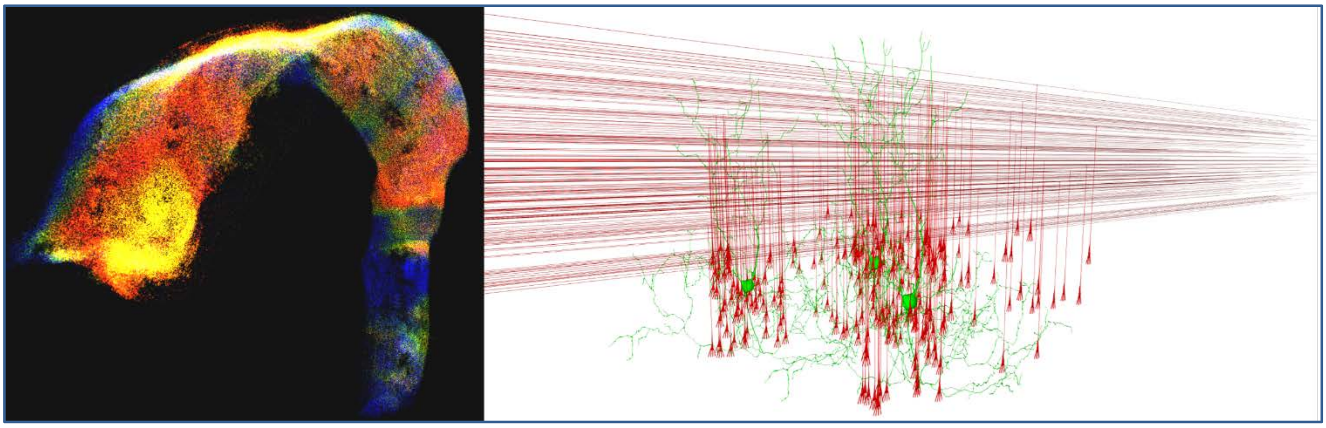

Figure 1: Visualisation of simulations and morphologies

The picture on the left side shows a model of the hippocampus, visualised with ViSimpl and simulated by Michele MIGLIORE (CNR; HBP Brain Simulation Platform) and his project partners using HPC resources at CINECA. The picture on the right side shows cerebellar granular layer and Golgi cells visualised with NeuroTessMesh (data provided by Egidio D'ANGELO's lab, Università di Pavia, Italy).

Project Number:	785907	Project Title:	Human Brain Project SGA2

Document Title:	SP7 High-Performance Analytics and Computing Platform – Results for SGA2 Year 2
Document Filename:	D7.6.2 (D47.2 D43) SGA2 M24 ACCEPTED 200731.docx
Deliverable Number:	SGA2 D7.6.2 (D47.2, D43)
Deliverable Type:	Report
Work Packages:	WP7.1, WP7.2, WP7.3, WP7.4, WP7.5, WP7.6
Key Result(s):	KR7.1, KR7.2, KR7.3, KR7.4; KR7.5
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA2 M24 / 31 Mar 2020
Actual Delivery Date:	SGA2 M26 / 06 MAY 2020; Accepted: 31 Jul 2020
Author(s):	Thomas LIPPERT, JUELICH (P20), SP Leader Maren FRINGS, JUELICH (P20), SP Manager
Compiled by:	Maren FRINGS, JUELICH (P20), SP Manager
Contributor(s):	Javier BARTOLOMÉ, Marta GARCIA GASULLA, Maria RIBERA SANCHO, David VICENTE, BSC (P5) Matthias BOLTEN, BUW (P7) Jacques-Charles LAFOUCRIERE, Gilles WIBER, CEA (P11) Michele CARPENE, Carlo CAVAZZONI, Giuseppe FIAMENI, Marco RORRO, Debora TESTI, CINECA (P13) Lucas BENEDICIC, Benjamin CUMMING, Colin MCMURTRIE, Alex UPTON, Stuart YATES, ETHZ (P18) Sandra DIAZ, Markus DIESMANN, Anne ELFGEN, Thorsten HATER, Wouter KLIJN, Sara KONRADI, Dorian KRAUSE, Charl LINSSEN, Abigail MORRISON, Boris ORTH, Alexander PEYSER, Dirk PLEITER, Jonas STAPMANNS, Bernd SCHULLER, Dennis TERHORST, Anna LUEHRS, Maren FRINGS, Anne NAHM, JUELICH (P20) Junxi WANG, Gabriel WITTUM, UFRA (P35) Hans Ekkehard PLESSER, NMBU (P44) Torsten KUHLEN, Simon OEHRL, RWTH (P46) Juan Pedro BRITO, Pablo TOHARIA RABASCO, UPM (P68) Susanna Mata FERNÁNDEZ, Luis PASTOR, Oscar David ROBLES, URJC (P69) Stefanie LACKNER, Sebastian RINKE, Hamid M. FARD, Felix WOLF, TUDA (P85)
SciTechCoord Review:	Mehdi SNENE, EPFL (P1)
Editorial Review:	Annemieke MICHELS, EPFL (P1)
Description in GA:	For consistent presentation of HBP results, SGA2 M24 Deliverables describing the accomplishments of an entire SP, WP or CDP have been prepared according to a standard template, which focuses on Key Results and the outputs that contribute to

	them. Project management elements such as Milestones and Risks will be covered, as per normal practice, in the SGA2 Project Periodic Report.
Abstract:	This Deliverable summarises the Key Results of the High-Performance Analytics and Computing (HPAC) Platform in the second year of SGA2. These comprise, among others, an updated architecture specification for the HPAC Platform and Fenix, data federation and data-intensive computing technology, Exascale-ready simulation technology, as well as interactive visual data analytics and *in situ* visualisation.
Keywords:	High Performance Analytics and Computing (HPAC) Platform, High Performance Computing (HPC), Platform, research infrastructure.
Target Users/Readers:	Computational neuroscience community, neuroimaging community, HPC community.

Table of Contents

Table of Tables

Table of Figures

1. Overview

The mission of the High Performance Analytics and Computing (HPAC) Platform is to build, integrate and operate the base infrastructure for the HBP, which comprises the hardware and software required to run large-scale, data-intensive, interactive simulations, to manage large amounts of data and to implement and manage complex workflows comprising concurrent simulation, data analysis and visualisation workloads.

As previously reported in SGA2 Deliverable D7.6.1 (D47.1 D22), a significant development was the start of the Interactive Computing E-Infrastructure (ICEI) project, funded through a separate Specific Grant Agreement (SGA) under the umbrella of the HBP Framework Partnership Agreement (FPA), which builds the Fenix infrastructure. Also in the second period of SGA2, the migration of HPAC Platform services to use the Fenix Infrastructure was a major work focus. The HPAC Platform also develops technologies for Fenix, e.g. the data transfer service, and it investigates technologies to support interactive data analysis and visualisation.

The HPAC Platform is engaged in improving HBP simulation software and ensuring their readiness for current and future HPC systems. Major updates of the NEST, Arbor and TVB simulators have been achieved in the past year. Various visualisation tools were further developed, made ready for being connected to running simulations through the *in-situ* pipeline, and the visualisation software catalogue is now available through the Knowledge Graph to broaden the scope and enhancing its usability.

2. Introduction

The mission of the High Performance Analytics and Computing Platform is to build, integrate and operate the base infrastructure for the HBP, which comprises the hardware and software required to run large-scale, data-intensive, interactive simulations, to manage large amounts of data, and to implement and manage complex workflows comprising concurrent simulation, data analysis and visualisation workloads.

A significant change in SGA2 was the start of the Interactive Computing e-Infrastructure (ICEI) project, funded through a separate Specific Grant Agreement (SGA) under the umbrella of the HBP Framework Partnership Agreement (FPA), which builds the Fenix Infrastructure. Fenix Infrastructure and HPAC Platform are tightly connected, as they are implemented and operated by the same five supercomputing centres. CEA of France, also a partner in the ICEI project, joined the HPAC Platform at the start of SGA2, and although the technical integration of its Très Grand Centre de Calcul (TGCC) HPC facility was significantly more complicated than expected, the new infrastructure components are now available. Neuroscience-specific platform services, which are running on top of the e-Infrastructure services, are designed, deployed and operated. Details of these services are being worked out taking the results of use case analysis into account. To validate architecture and implementation details, a set of validation tests have been formulated and implemented. To further exploit the generated knowledge, efforts have been made to transfer this knowledge to HPC and Cloud solution providers, to improve their understanding of the needs of the HBP and the future EBRAINS Research Infrastructure (see Key Result KR7.1). Compute and storage resources are provided via the ICEI resource allocation mechanisms. In the last few months, progress has been made in the base infrastructure and infrastructure services for the HBP Platforms in a number of areas, including the development of the container runtime engine Sarus tailored to HPC systems, continued support for the Neurorobotics Platform, and also increased collaboration with SP5 (see Key Result KR7.5).

In the past year we have seen an increasing amount in HBP research groups using Fenix IT infrastructure resources; the required implementation support being provided by the HPAC Platform. Scientists outside the HBP can also apply to use Fenix resources via PRACE and hence use the same infrastructure as HBP scientists. The HPAC Platform also develops technologies for Fenix (see Key Result KR7.2), e.g. the data transfer services, that have been enabled in order to allow HBP users to move datasets between all five HPC sites.

Significant efforts were also invested in the improvement of simulation software and work to ensure their readiness for current and future HPC systems (see Key Result KR7.3). The capabilities of the NEST simulator were considerably advanced with releases 2.18.0 and 2.20.0. The release of NEST 3.0 with a much more expressive and efficient interface for network construction is forthcoming. Voucher-funded work on the NEST Desktop graphical user interface will make advanced computational modelling accessible to a wider audience and become a powerful tool for neuroscience education. Leading international groups are actively using NEST today and contributing back to it and thus HBP, e.g. the Allen Brain Institute, the Okinawa Institute of Technology and RIKEN. Further work focused on reproducibility and standardisation. Great efforts have been invested in the release of a first version of TVB-HPC which connects neural mass modelling with high performance computing in order to allow highly efficient parameter fitting. For Arbor, which has been developed entirely within the HBP to provide efficient, state-of-the art, future proof simulation technology, a first full-featured and user accessible version has been released within the last year.

Other important work undertaken within the HPAC Platform concerns interactive visual data analytics and in-situ visualisation tools and techniques (see Key Result KR7.4). Various visualisation tools were further developed, made ready for being connected to running simulations through the *in-situ* pipeline, and the visualisation software catalogue is now available through the Knowledge Graph (KG) to broaden the scope and enhancing its usability.

To ensure that users can get the most out of the HPAC Platform, the HPAC Guidebook[1], which collects all relevant documentation, has been continuously updated. Dissemination, education and training

[1] https://hbp-hpc-platform.fz-juelich.de/?page_id=2215%20

events have been organised, often as joint efforts between ICEI and HPAC. The HPAC team has supported scientists in their work through the central support team and, at a more advanced level, the HLST and the performance optimisation service.

3. Key Result KR7.1 Architecture specification for the HPAC Platform and Fenix

One of the key aims of the HBP is to build a Research Infrastructure (RI) to support research and development by the neuroscience communities. This RI strongly relies on e-infrastructure services to enable neuroscientists, e.g. to

- Deploy community-specific Platform services that are typically web-based and are in the future branded as EBRAINS services
- Facilitate curation of data
- Run large-scale, data intensive, interactive multi-scale brain simulations (up to the scale of a full human brain)
- Manage the large amounts of data used and produced by simulations and in experiments
- Manage complex workflows comprising concurrent simulation, data analysis and visualisation workloads

The generic e-infrastructure services including HPC, Cloud and data services are being provided by the ICEI project. The portfolio of ICEI services comprises of Scalable Computing Services, Interactive Computing Services, Virtual Machine Services, Active Data Repositories and Archival Data Repositories. These services are integrated through a relatively thin layer of federation services, including Authentication and Authorization Infrastructure (AAI) services and the Fenix User and Resource Management (FURMS) services. Enabling a flexible composition of these e-infrastructure services is key to adjust to the changing needs of the EBRAINS RI.

Using Fenix e-infrastructure services, the HBP is deploying a variety of Platform services that are largely specific for the neuroscience community. During SGA2, such services have been developed in different Subprojects. The HPAC Platform specification focusses on a sub-set of services that are crucial to enable complex workflows that exploit high-end computing, networking and storage resources in a distributed and federated e-infrastructure. The following services have been architected and documented in deliverable SGA2 Deliverable D7.1.1 (D42.1 D81):

- Software deployment services (SWD): These services facilitate the availability of HBP software products across the Fenix IT infrastructure
- Developer services (DEV): Services needed for software development including standard source code management systems, bug tracking systems and continuous integration systems
- Database services (DB): Structured and unstructured database (DB) and data base management system (DBMS) used for storing HBP/EBRAINS data
- Data management (DMGT): Interfaces to information about data stored in Fenix Archival Data Repositories and their transfer and federation for different Platform services
- Monitoring services for Platforms (MONP): Tools and services that allow Platform users and the Platforms themselves to access information on the availability of various Platform services
- Visualisation services (VIZ): Application-specific and HPC visualisation services for end-users and Platform developers

Providing a more specific design description of each of these services is in progress. This work is performed in close conjunction with an analysis of use cases.

3.1 Outputs

3.1.1 Overview of Outputs

3.1.1.1 List of Outputs contributing to this KR

- Output 1: HPAC/Fenix SGA2 architecture document v01 and v02 (C2669, C2670): Architecture document on the HPAC Platform services and their relation to ICEI e-infrastructure services (see Table 1)
- Output 2: SP7 Use Case and requirements documentation (C2671): Final list of SP7 science and Use Cases augmented with information about needs for HPAC Platform services (see Table 2).
- Output 3: HPAC SGA2 Platform validation report (C2673): Final report on validation tests that had been defined and those that have been executed during SGA2 (see Table 3).

3.1.1.2 How Outputs relate to each other and the Key Result

The use and science cases (Output 2, C2671) have been a key for defining the initial architecture as well as the refined definition of the different services (Output 1, C2669 and C2670). The validation tests listed in Output 3 (C2673) allow to verify whether the realised services, e.g. those defined in the architecture document, do meet the initially defined needs.

3.1.2 Output 1: HPAC/FENIX SGA2 architecture document

Table 1: KR7.1 Output 1 links

Component	Link to	URL
C2669	User Documentationy	see SGA2 Deliverable D7.6.1 (D47.1 D22)
C2670	User Documentation	see SGA2 Deliverable D7.1.1 (D42.1 D81)

This Output provides an architecture for the High Performance Analytics and Computing (HPAC) Platform. It describes a set of Platform services that are being realised on top of the e-infrastructure services provided by the ICEI project involving five leading supercomputing data centres in Europe. The initial specification was provided as Deliverable D7.1.1 (referred to as "HPAC/FENIX SGA2 architecture document v1"). Later a complementary document was produced (referred to as "HPAC/FENIX SGA2 architecture document v2") that assesses the current status of the different services, the connections to other documents produced by SP7 during SGA2 as well as an initial list of Platform service specifications including recommendations for next steps in SGA3.

3.1.3 Output 2: SP7 use case and requirements documentation

Table 2: KR7.1 Output 2 links

Component	Link to	URL
C2671	Data Repository	https://wiki.humanbrainproject.eu/bin/view/Collabs/hbp-ebrains-use-cases

Set of science and Use Cases that has been collected throughout SGA2 as well as within the ICEI project. In this context, the documentation of the science and use cases has been augmented with an analysis of the specific needs for services described in component C2669 “HPAC/Fenix SGA2 architecture document v1”.

3.1.4 Output 3: HPAC SGA2 Platform validation report

Table 3: KR7.1 Output 3 links

Component	Link to	URL
C2673	Data Repository	https://collab.humanbrainproject.eu/#/collab/80671/nav/546311?state=uuid%3D560ad356-d5a8-408e-a4e7-0bded11e6606

Collection of test definitions aiming for a validation of the HPAC Platform as well as a documentation of the test results collected by the end of SGA2.

3.2 Validation and Impact

3.2.1 Actual and Potential Use of Output(s)

- The architectural design provided with Component C2669 will be used to further develop the HPAC services in the context of EBRAINS.
- The science and use cases provided with Component C2671 continue to be a key basis for further co-design of EBRAINS services.
- The validation tests provided with Component C2673 will continue to be used to assess future versions of the Platform.

3.2.2 Publications

No publications yet

4. Key Result KR7.2 Data federation and data-intensive computing technology

This phase of SGA2 focussed on finalising the data location and transfer services, providing user-level documentation, and documenting the experiences made in SGA2. Together, these results provide a data federation layer for the HBP, leveraging infrastructure provided by the Fenix project, as well as mature technologies developed by HBP member sites. In particular, these are the Knowledge Graph, UNICORE, and the Fenix Archival Data Repositories.

For these services technologies were chosen that are already in production inside the HBP, which enables synergies between different teams and user communities.

Further, research into novel storage technologies and interactive data analysis has continued. This research has been made accessible to users and developers inside the HBP in a form suitable to enable them to apply our results to their projects. This has been done in two parts, one being documentation and educational material, the second the release of higher-level abstractions over these technologies, namely extensions to the SLURM resource manager and the Hecuba interface for distributed Key-Value Stores.

4.1 Outputs

4.1.1 Overview of Outputs

4.1.1.1 List of Outputs contributing to this KR

- Output 1: HBP/Fenix data transfer service (C416, see Table 4)

  Service to enable users to move data between locations internal and external to the HBP.
- Output 2: HBP/Fenix data location service (C2826, see Table 5)

  Service to track, index, and search data collections relevant to the HBP.
- Output 3: Private Cloud On a Compute Cluster (PCOCC, C2825, see Table 6)

  Software package to emulate cloud computing on a HPC cluster.
- Output 4: Data Access and Efficient I/O (C2826, see Table 7)

  A knowledge base for users to facilitate education in I/O technologies for the HBP.
- Output 5: SLURM co-scheduling plugin (C2824)

  An extension to the SLURM resource manager for efficient use of high-performance storage tiers.

4.1.1.2 How Outputs relate to each other and the Key Result

The services of data location and data transfer services (Outputs 1 and 2) are expected to be tightly integrated, as the indexing and tracking of metadata directly feeds into obtaining, replicating, and moving the actual data items. This is not complete yet and will grow further based on user feedback and requirements.

Output 3, PCOCC, provides a way for users to use cloud and HPC resources transparently. Both types will directly use data location and transfer services.

In Output 4, education for end users and HBP developers is provided to enable them in the use and optimisation of storage solutions offered by the HBP partners and Fenix. A particular focus lies on the use of node-local, high-performance storage (NVMe). Further, experience reports of application owners are included. These documents summarise many of the experiences made by us and our partners in the course of SGA2.

The SLURM co-scheduling plugin (Output 5) allows users to handle modern I/O capabilities in a more abstract fashion, while sharing these resources with others. PCOCC is expected to benefit from the use of the co-scheduling plugin.

4.1.2 Output 1: HBP/Fenix data transfer services enabled

Table 4: KR7.2 Output 1 links

Component	Link to	URL
C416	Technical Documentation	https://www.unicore.eu/docstore/uftp-2.2.0/uftp-manual.html
	User Documentation	https://collab.humanbrainproject.eu/#/collab/3656/nav/275433

The Fenix data transfer services (C416) have been enabled in order to allow HBP users to move datasets between all the HPC sites, both in "client to server" and "server to server" mode (see Figure 2). The technology adopted in HBP to implement the Fenix data transfer service is the UNICORE UFTP. The UFTP protocol is a high-performance data transfer service based on FTP, developed as part of the UNICORE suite.

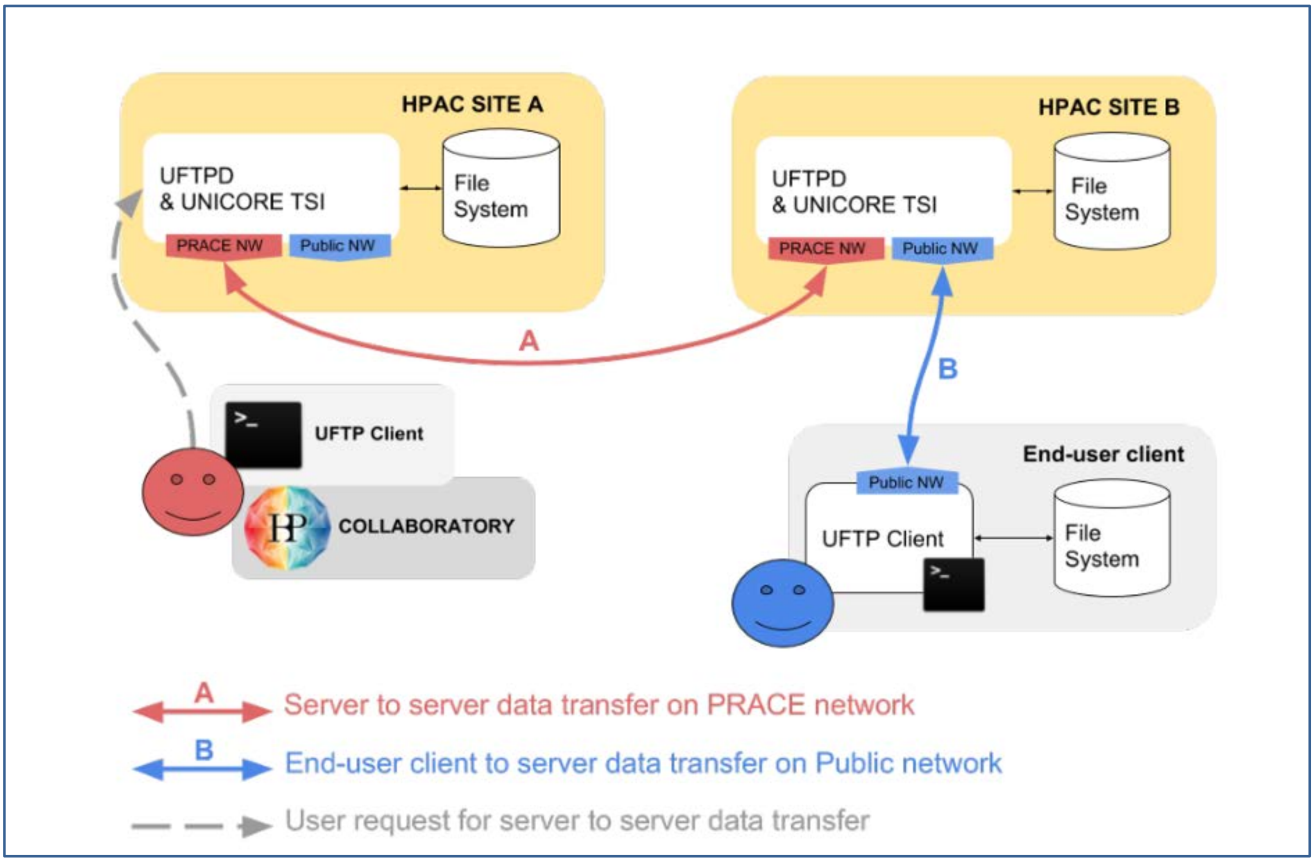

Figure 2: General view of the data transfer services in the HPAC Platform and Fenix

This figure shows the main components involved in data transfer between the sites of the HPAC Platform and the Fenix infrastructure. Server to server and client to server connections are highlighted in red and blue, respectively.

HBP users are able to use the UFTP service to move files using two different interfaces:

- The UFTP standalone client

  A fully featured client for the UNICORE middleware, able to start both direct data transfer (user workstation to server) and third party transfers.

- The UNICORE REST API

  Provides an API for driving UNICORE and UFTP via HTTP verbs, which can be used from IPython Notebooks, and, thus, the Collaboratory. The pyunicore library provides an abstraction over the REST API suitable for end users.

4.1.3 Output 2: HBP/Fenix data location service (C2826)

Table 5: KR7.2 Output 2 links

Component	Link to	URL
C2826	Technical Documentation	https://bbp-nexus.epfl.ch/staging/docs/
	User Documentation	https://bluebrainnexus.io/

The HBP Central Data Location service (C2826) provides functionality for tracking the location of data across the HBP, discovery and search of such collections, tracing of versions, and — to a limited extent – provenances to HBP users. It leverages the HBP Knowledge Graph (KG) service developed by EPFL, which offers a production-ready service for searching and linking entries loosely adhering to a collection of schemata. Using the KG as the basis of the data location service also has the benefit of automatically providing unique, persistent IDs for each data collection.

The Knowledge Graph (KG)

- KG Search UI, Anonymously accessible:

https://www.humanbrainproject.eu/explore-the-brain/search

- KG Editor & Query Builder, Protected Access:

  https://kg.humanbrainproject.eu/editor
- KG Query API, Protected Access:

  https://kg.humanbrainproject.org/apidoc/index.html?url=/apispec/spring%3Fgroup%3D00_external

For access to protected/private sections, an HBP login is required, access permission can be requested at kg-team@humanbrainproject.eu.

In the near future, the integration between the HBP Central Data Location Service (CDLS) and the Fenix/HBP Archival Data Repositories (ARD) (based on OpenStack SWIFT) will be completed. This step is needed to allow HBP workflows to leverage Fenix ARD for long term storage of results cross-referenced in the CDLS. To this end, we will provide a collection of metadata schemata suited for tracking locations, types, and provenance of such data.

We demonstrate how such workflows can be built in the current state of affairs regarding authentication and federation in the "KG APIs + SWIFT integration" demo example notebook, which can be accessed by HBP users under:

https://collab.humanbrainproject.eu/#/collab/49291/nav/338055

Further efforts in the direction of support libraries, mainly in the Python language, will be investigated to enable end users to leverage these services.

Further, efforts have been made to export HBP metadata from the KG via standard protocols (OAI-PMH to an external B2FIND service such that public KG data can be discovered via the EOSC portal.[2] The documentation for the KG harvester is available at https://gitlab.version.fz-juelich.de/hater1/reaper.

4.1.4 Output 3: Private Cloud On a Compute Cluster (PCOCC)

Table 6: KR7.2 Output 3 links

Component	Link to	URL
C2825	User Documentation	https://pcocc.readthedocs.io
	Technical Documentation	https://github.com/cea-hpc/pcocc

PCOCC (pronounced "peacock": Private Cloud on a Compute Cluster) allows HPC cluster users to host their own clusters of Virtual Machines (VMs) on compute nodes alongside regular HPC jobs through SLURM. Such VMs allow users to fully customise their software environments for development and testing, or for facilitating application deployment. Compute nodes remain managed by the batch scheduler, as usual, since the clusters of VMs are treated as regular jobs. From the point of view of the batch scheduler, each VM is a task for which it allocates the requested CPUs and memory, and the resource usage is accounted to the user, just as for any other job. For each virtual cluster, PCOCC instantiates private networks isolated from the host networks, creates temporary disk images from the selected templates (using Copy-on-write) and instantiates the requested VMs. PCOCC is able to run virtual clusters of several thousands of VMs and will enable different usage scenarios for compute clusters, from running complex software stacks packaged in an image, to reproducing technical issues happening at large scale without impacting production servers. PCOCC has been developed at CEA. A PCOCC version is in production, which can now support virtual machines and containers in an HPC context. This tool provides to users a single, "easy to use" interface for running a virtual HPC cluster with very good performance. The description of this tool (the full documentation is on https://pcocc.readthedocs.io)

[2] At the time of writing of this report the implemented harvester replicates KG data into a test instance. For an example see: http://eudat7-ingest.dkrz.de/dataset/f6ffa725-2348-53bc-96a5-8119ef7f2d44

4.1.5 Output 4: Data Access and Efficient IO

Table 7: KR7.2 Output 4 links

Component	Link to	URL
C2823	Technical Documentation	https://wiki.humanbrainproject.eu/bin/view/Collabs/how-to-data-access-and-efficient-io

Task T7.2.3 has developed a series of guidelines, examples, and user-level documentation describing how to make efficient use of modern I/O technologies.

Topics covered are the following

- Using the SWIFT object storage for multi-site workflows. Educates users on the use of federated data services offered by Fenix. An experience report of the "learning to learn" project team is included, describing the construction of a federated workflow.
- Leveraging modern MPI for reducing I/O load. Demonstration of how to make best use of I/O bandwidth.
- Best practices for using HDF5 for image data. Covers the use of HDF5 in ML/DL, an important use case for the BrainAtlas among others.
- Using BeeGFS-on-demand for high-performance I/O across node-local NVMe devices. Demonstrates use of an ad hoc high-performance file system for data intensive computing.
- Usage and performance of Apache Pass Optane devices. Leveraging novel memory technologies for persistent, memory intensive workloads.
- Interactive visualisation using the Universal Data Junction. Proof of concept for coupled, lock-step simulation and visualisation.
- Hecuba as an abstraction for distributed Key Value Stores (KVS). Use Case for a high-level library leveraging industry-strength KVS for HPC.

This collection is publicly available in the Collaboratory.

https://wiki.humanbrainproject.eu/bin/view/Collabs/how-to-data-access-and-efficient-io.

4.1.6 Output 5: SLURM Co-Scheduling Plugin

Table 8: KR7.2 Output 5 links

Component	Link to	URL
C2824	Technical Documentation	https://collab.humanbrainproject.eu/#/collab/8122/nav/61636?state=uuid%3D7f322b30-49f4-473d-8e27-40c42976d1d5
	User Documentation	https://github.com/HumanBrainProject/coallocation-slurm-plugin/blob/master/README.md

Work on the development of a plug-in solution for SLURM (C2824), which couples compute and high-end storage resources such as SSD/NVM storage volumes, and adds the ability to schedule both types of resources in a way that balances system utilisation and turn-around time, was delayed due to unforeseeable circumstances. In the last months, a basic version of the SLURM plugin has been implemented and tested in a virtual cluster and is published under the GNU General Public License (GPLv3). The plugin has been tested in a virtual cluster but is not yet integrated in the HPAC Platform.

4.2 Validation and Impact

The UNICORE FTP (UFTP) service federation allows HBP users to transfer data between the Fenix sites (BSC, CEA, CINECA, ETHZ-CSCS and JUELICH-JSC) and KIT in a secure and seamless manner, thus improving the capabilities of the infrastructure that the HPAC Platform builds on. The viability and usability of these solutions has been demonstrated.

PCOCC enables new utilisation models of HPC resources, as users will be able to run virtual clusters of several thousands of VMs. The SLURM co-scheduling plugin will enable shared use of integrated NVMe storage for HPC clusters.

4.2.1 Actual and Potential Use of Output(s)

The data transfer services have already been successfully exploited by the HBP Co-Design Project CDP2 “Mouse-Based Cellular Cortical and Sub-Cortical Microcircuit Models” to transfer output data sets between CINECA and JUELICH-JSC.

The introduction of the PCOCC also reinforced this model and helped users to execute deeply customised environments, using HPC resources through virtualisation.

The federation capabilities of the archival data repositories have been exploited in the learning to learn framework, as reported in Output 4.

4.2.2 Publications

No publications yet.

5. Key Result KR7.3 Exascale-ready simulation technology: NEST, Arbor and TVB

The focus of KR7.3 is the improvement of simulation software and work to ensure their readiness for current and future HPC systems. We have considerably advanced the capabilities of the NEST simulator with releases 2.18.0 and 2.20.0 and are very close to releasing NEST 3.0 with a much more expressive and efficient interface for network construction.

During the second phase of SGA2 we further worked towards reproducibility and standardisation. The development and refinement of domain-specific languages such as NESTML (to support the automatic code generation of neuron and synaptic models at point neuron scale) and NeuroML (to support automatic code generation of neural mass models targeting multiple back ends such as CUDA and numba) help neuroscientists use simulation engines in an easier and more efficient way, while preserving the flexibility required to allow new models to be incorporated by end users.

Progress concerning TVB-HPC was the release of a first version, including parameter-fitting framework and automatic code generation of simulation kernels. This Output connects neural mass modelling with high performance computing in order to allow highly efficient parameter fitting.

For Arbor, the simulation software for large networks of multi-compartment cells on HPC systems, written as a C++ library with a Python front end. Work over M13-M24 has focused on the Arbor core C++ library and its Python wrapper, to release the first full-featured and user-accessible version of Arbor.

Moreover, the The NEST user experience and sustainability has improved strongly. Based on a workshop on documentation approaches early in SGA2, NEST documentation has been re-worked technically (now available on readthedocs.io[3]) and revised for completeness, correctness and consistency. At the same time, the NEST community has been further strengthened with very successful NEST conferences in 2018 and 2019 and regular NEST hackathons bringing developers and in some cases users together for focused development sprints. Outside the conference and hackathons, NEST developers meet every other Monday for an Open Developer Video Conference to discuss current topics in development and review open pull requests, issues and mailing list requests.

[3] https://nest-simulator.readthedocs.io/en/stable/

Between 25 and 40 developers have contributed code, documentation or examples to recent NEST releases. Systematic code review, continuous integration testing and refactoring maintain and improve NEST code quality ensuring long-term sustainability.

5.1 Outputs

5.1.1 Overview of Outputs

5.1.1.1 List of Outputs contributing to this KR

- Output 1: NEST (incl. NESTML) (C209, C2696, C510, C2679, C2678, C2789, C661, C662, see Table 8)
- Output 2: Arbor (C2689, C2690, C2691, see Table 9)
- Output 3: TVB (C3030, see Table 10)
- Output 4: Multi-simulator multi-scale interaction (C2692, C2693, C2694, C2695, see Table 11)

5.1.1.2 How Outputs relate to each other and the Key Result

The Arbor, NEST and TVB Outputs all perform simulation of brain activity, however at different levels of detail. Arbor simulates networks of highly detailed cells with complicated dynamics, NEST can simulate much larger networks by using simpler models for individual cells, and TVB can scale up to models of the entire brain by simulating regions in place of individual cells. Together, these Outputs provide users of the HPAC Platform state-of-the-art simulation tools for a broad range of simulations.

The Output TVB-HPC describes the extensions of TVB in order to optimally utilise the HPC infrastructure provided within the HBP. This Output is no prerequisite for any other Output but moves the infrastructure development cohesively together with other simulation Components developed in SP7 which leverage high performance computing while providing a usable interface for model definition.

These Outputs are complemented by implementations of key features extending the capabilities of the tools and simulators (IDs: C510, C2690, C2692, C2693, C2694, C2696), as well as enhancements in usability and maintainability (ID: C2678, C2679). In addition to user documentation, active support of user communities has been established (IDs: C661, C662, C2689, C2789).

5.1.2 Output 1: NEST (incl. NESTML)

Table 9: KR7.3 Output 1 links

Component	Link to	URL
C209	Software Repository	https://github.com/nest/nest-simulator
	Technical Documentation	https://nest-simulator.readthedocs.io/en/stable/
	User Documentation	https://www.nest-simulator.org/documentation/
C510	Software Repository	https://github.com/nest/nest-simulator
	Technical Documentation	https://github.com/ReScience-Archives/Senden-Schuecker-Hahne-Diesmann-Goebel-2018

	User Documentation	https://github.com/ReScience-Archives/Senden-Schuecker-Hahne-Diesmann-Goebel-2018
C2679	Software Repository	https://github.com/nest/nestml
	Technical Documentation	https://github.com/nest/nestml/pull/459/files#diff-1e77cba9d6acc33693cda56a75663f9d
	User Documentation	https://github.com/nest/nestml/blob/036b07469b2df65a3d04dcbb527e419374ffef10/doc/synapses_in_nestml.rst
C2789	Software Repository	https://github.com/nest/nest-simulator
	Technical Documentation	https://collab.humanbrainproject.eu/#/collab/264/nav/325519
	User Documentation	https://nest-simulator.org/documentation
C2696	Software Repository	https://github.com/nest/nest-simulator
	Technical Documentation	https://github.com/nest/nest-simulator/pull/1095 and https://github.com/nest/nest-simulator/pull/1257
	User Documentation	https://nest-simulator.readthedocs.io/en/v2.20.0/models/clopath_synapse.html
C2678	Software Repository	https://github.com/nest/nest-simulator

In the second year of SGA2, NEST 2.18.0 (M15) and NEST 2.20.0 (M22) were released (C209). At the same time, online documentation for NEST has been strongly improved, both in technical delivery (now on readthedocs.io[3]), consistency and quality. NEST now incorporates the infrastructure for the representation of third factor plasticity rules. As an example, the Clopath plasticity rule was released with NEST 2.18.0. In addition, a pull request against NEST has been made for the compartmentalised third factor plasticity rule by Urbanczik and Senn (C2696). The prototype implementation of a vastly improved, unified user interface and back end for connectivity generation (C2678) was integrated into the NEST master branch early in M23 and is thus available for public testing; at this time, it is planned to release this work as part of NEST 3.0 towards the end of M24.

A new NESTML version has been released (C2679), which improves the development workflow and usability for the end user and enhances the automated processing of differential equations.

Functional validation of neuron model correctness has been implemented, as an important step towards improving the NEST Simulator code base maintainability in the coming years — see also section Conclusion and Outlook below.

The specification of synapse models has been enhanced to cover advanced plasticity rules including spike-timing dependent plasticity (STDP).

5.1.3 Output 2: Arbor

Table 10: KR7.3 Output 2 links

Component	Link to	URL
C2690	Software Repository	https://github.com/arbor-sim/nsuite
	Technical Documentation	https://nsuite.readthedocs.io
	User Documentation	https://nsuite.readthedocs.io
C2691	Software Repository	https://github.com/arbor-sim/arbor
	Technical Documentation	https://arbor.readthedocs.io/
	User Documentation	https://arbor.readthedocs.io/

Arbor is simulation software for large networks of multi-compartment cells on HPC systems, written as a C++ library with a Python front end. It has been developed entirely within the HBP to provide efficient, state-of-the-art, future-proof simulation technology. Arbor will facilitate beyond state-of-the-art modelling in the HBP in three key ways. First, Arbor is the only multi-compartment simulation engine to natively support NVIDIA GPUs, and is also the only that can be easily extended to new architectures. Arbor also provides a model description interface that is both flexible and user-friendly, with a clean separation between model description and implementation, required for portable model descriptions. Finally, Arbor provides a rich API for input and output, which will facilitate robust coupling with other simulation tools (namely NEST and TVB on EBRAINS in SGA3).

Work over M13-M24 has focused on the Arbor core C++ library and its Python wrapper (C2691), to release the first full-featured and user-accessible version of Arbor.

The Python wrapper for Arbor was released with full documentation, including a "getting started" guide, available online[4].The wrapper makes all features of the C++ library available to scientists, plus a Python-only single cell workflow that facilitates rapid model development and testing. It has been tested on laptops, desktop computers and HPC clusters at CSCS, with support for all CPU and GPU architectures available through the HPAC Platform.

One of the most challenging and time-consuming aspects of modelling morphologically-detailed cells is describing the morphology and the location of ion channels, synapses and electrical properties thereon. This was addressed by Arbor's new "morphology engine", for describing cell models and properties that are both flexible and easy for users to read and write:

- a low-level morphology representation compatible with morphologies from SWC, Neurolucida and NeuroML file formats;
- a Domain Specific Language (DSL) for compactly representing regions and sets of locations on a morphology;
- a method for associating ion channels and synapses with regions and locations;
- a new back end that hides implementation details from end users.

Support for the NMODL language, used to describe ion channel and synapse dynamics, has been improved significantly to be able to support all widely-used dynamics.

The NSuite benchmarking and validation suite (C2690; released previously in SGA2) has been updated to benchmark and validate the latest version of Arbor.

5.1.4 Output 3: TVB

Table 11: KR7.3 Output 3 links

Component	Link to	URL
C3030	Software Repository	https://github.com/the-virtual-brain/tvb-hpc
	Technical Documentation	https://github.com/the-virtual-brain/tvb-hpc
	User Documentation	https://github.com/the-virtual-brain/tvb-hpc

During the second year of SGA2 we released the first version of TVB-HPC including the parameter-fitting framework and automatic code generation of simulation kernels. The current release is available via the public GitHub repository (https://github.com/the-virtual-brain/tvb-hpc). The current version of TVB-HPC includes automatic code generation for CUDA, Python and numba kernels generated from a NeuroML description. NeuroML, together with LEMS, has been extended in order to support the requirements of neural mass models. The first set of unit tests have been added to the release and the full benchmarking suite will be available by the end of M24. An example workflow to generate automatic kernels from the Collaboratory will be also available by M24.

[4] https://arbor.readthedocs.io/en/latest/python.html

This Output connects neural mass modelling with high performance computing in order to allow highly efficient parameter fitting. This will substantially reduce the time used to bring models closer to experimental data and will allow for a better understanding of the mechanisms of the brain at this scale of description. Automatic code generation is also essential to lower the entry barrier for researchers around the field of neuroscience to efficiently using simulation technology.

5.1.5 Output 4: Multi-simulator multi-scale interaction

Table 12: KR7.3 Output 4 links

Component	Link to	URL
C2692	Software Repository	https://gcsc.uni-frankfurt.de/simulation-and-modelling/projects
	Technical Documentation	https://durus.gcsc.uni-frankfurt.de/~neugen/publications.html
	User Documentation	https://durus.gcsc.uni-frankfurt.de/~neugen/documentationneugenmanual.html
C2693	Software Repository	Music- and Neuron-UG-plugins in UG4 (https://github.com/UG4/ugcore)
	Technical Documentation	N/A (not yet released) ug4_music_specification.pdf (see https://collab.humanbrainproject.eu/#/collab/264/nav/1975 nest_neuron-ug.md
C2694	Model Repository	http://doi.org/10.3389/fncom.2015.00094
C2695	Technical Documentation	http://ug4.github.io/docs/
	User Documentation	https://github.com/UG4/ugcore/blob/master/README.md

The first version of the NeuGen-VRL-Plugin has been published as planned (C2692). The users can directly generate and visualise the neuron networks in VRL-Studio. To achieve this, NeuGen was optimised to minimise its dependence on the GUI. New VRL interfaces were implemented, and visualisation functions had to be newly implemented using VRL libraries.

UG4 can receive and send continuous data through the music-plugin, thus it can receive the membrane potential sent by NEURON using the Neuron-UG-plugin (C2693). To achieve this, Music-dependent classes were implemented in UG4, and a MusicRuntime Object was created in LUA-script for calling the tick function and accessing the time provided by MUSIC. More details can be found in ug4_music_specification.pdf[5].

A model with full spatial resolution for the extracellular potential activity was developed using ug4 (C2694). The creation of a model which uses transmembrane currents (for example delivered by MUSIC) to calculate the extracellular potential in three-dimensional space is documented in a PhD Thesis which is not published yet. See ug4_music_specification.pdf[5] for further details.

C2695 (software that couples the relevant scales for network simulation with a 3D model) is still being processed since it strongly depends on Component ID 2693 (NEST-NEURON-UG Coupling) and Component ID 2694 (3D model).

5.2 Validation and Impact

5.2.1 Actual and Potential Use of Output(s)

The simulators NEST, Arbor and TVB are an essential part of the HBP software infrastructure and enable brain research spanning different scales of resolution.

Output 1 (NEST):

[5] https://collab.humanbrainproject.eu/#/collab/264/nav/1975?state=uuid%3D2c78cd7d-1e80-41ba-8312-2663240bcbf8

NEST has been widely adopted within and outside the HBP. Use-case driven development and active support of users has strengthened and extended the user community and also encouraged community contributions.

Numerous groups across science and infrastructure HBP Subprojects use NEST, including SP3 (PENNARTZ, Amsterdam; PAOLUCCI, Rome; STORM, Oslo), SP4 (SENN, Bern; EINEVOLL, Ås; DIESMANN, Jülich), SP10 (Neurobotics Platform), CDP2 (D'ANGELO, Pavia), CDP4 (GOEBEL, Maastricht), CDP5 (PETROVICI, Bern) and Partnering Project CerebNEST (PEDROCCHI, Milano). Close interaction with these user groups guides NEST development by providing essential use cases. In collaboration with SP9 (FURBER, Manchester), we have also established NEST as the gold standard for validating neuromorphic computing systems. By addition of the NEST Desktop through a Voucher-funded Partnering Project, NEST will become available as a powerful education tool, ensuring further uptake.

Leading international laboratories use NEST in their brain modelling efforts: The Allen Brain Institute, who use NEST to implement large parts of their mouse visual system model (MIHALAS); this group has contributed their GLIF neuron model class to NEST 2.20.0, integrating high-quality code in collaboration with the HBP High-Level Support Team (HLST). The DOYA-group at the Okinawa Institute of Science and Technology uses NEST as the simulation engine for their very-large-scale models of the primate brain, closely collaborating with the RIKEN Center for Computational Science (IGARASHI) on model scaling on pre-exascale systems.

Output 2 (Arbor):

Arbor has been developed from scratch during SGA1 and SGA2, as such it has not had enough feature support to generate scientific results to date. A detailed Purkinje cell model from component C1646, developed by SGA2 Task T6.2.2 at University of Pavia, has been used to motivate development of Arbor over the last 12 months. The Purkinje cell has been implemented fully in Arbor's Python wrapper, and collaboration will continue with University of Pavia to support the other cell types from C1646.

Output 3 (TVB):

The JIRSA-group in Aix-Marseille and the INM-7 at Jülich (EICKHOFF) are active users of the TVB-HPC DSL code. In particular, Thanos MANOS, who was at INM-7 but is now working at the University CY Cergy in Paris, has used and continues to use this code. He is using it for fast and extensive parameter fitting of neural mass models on GPUs, with the goal of understanding the effect of using different atlases on the ability of mathematical models to reproduce experimental functional and structural connectivity data.

5.2.2 Publications

- Akar et al., Arbor — A Morphologically-Detailed Neural Network Simulation Library for Contemporary High-Performance Computing Architectures (P1789)

This publication introduces Arbor, a performance portable library for simulation of large networks of multi-compartment neurons on HPC systems

- Einevoll et al., The Scientific Case for brain simulations (P1913)

This publication argues why simulators will likely be indispensable for bridging the scales between the neuron and system levels in the brain, and why a set of brain simulators based on neuron models at different levels of biological detail should therefore be developed.

6. Key Result KR7.4 Interactive visual data analytics and *in situ* visualisation developed and deployed on Fenix and HPAC Platform infrastructure

The focus in the second period of SGA2 was on the continuation of software development for *in situ* visual analysis of simulation results, including a redesign of the in situ pipeline that connects simulation and analysis tools, and on the implementation of a back end that enables visualisation and analysis tools to run on the new Fenix infrastructure. The back end has been successfully deployed at CINECA (P13) in a close co-development effort with Michele MIGLIORE (CNR, P12) and his group. The visualisation software catalogue has been released through the Knowledge Graph of the HBP and additionally includes metadata on general software relevant for HBP and EBRAINS. The latter has been developed in close collaboration with the Knowledge Graph, Collaboratory, and data curation teams.

6.1 Outputs

6.1.1 Overview of Outputs

6.1.1.1 List of Outputs contributing to this KR

- Output 1: In situ Pipeline (C2680, C2697, see Table 12)
- Output 2: Deployment and extension of visualisation tools at CINECA (C2681)
- Output 3: Release of a prototype for remote visualisation and interaction (C2682, see Table 13)
- Output 4: Visualisation software catalogue – design and first version (C2707, see Table 14)

6.1.1.2 How Outputs relate to each other and the Key Result

Outputs 1, 2, and 3 are related since Output 3 is planned to allow users to establish visualisation sessions on the HBP infrastructure as well as to manage sessions for running simulators within *in-situ* visualisation scenarios (Output 1). Output 2 will ease the deployment of the visualisation tools on the supercomputers within the HBP infrastructure.

Output 4 presents an entry point to software for visualisation and, thus, intends to represent an element for community building. The latter is also reflected by bi-weekly video conferences on visualisation tools and overall support provided on request. In the second year of SGA2, Michael DENKER (JUELICH, P20) received support for the development of visualisation of higher order correlations.

6.1.2 Output 1: In situ Pipeline

Table 13: KR7.4 Output 1 links

Component	Link to	URL
C2697	Software Repository	https://github.com/gmrvvis/visimpl
	Technical Documentation	https://github.com/gmrvvis/visimpl
	User Documentation	https://github.com/gmrvvis/visimpl
C2680	Software Repository	https://devhub.vr.rwth-aachen.de/VR-Group/in-situ-pipeline

	Technical Documentation	https://devhub.vr.rwth-aachen.de/VR-Group/in-situ-pipeline/insite/tree/develop
	User Documentation	https://devhub.vr.rwth-aachen.de/VR-Group/in-situ-pipeline/insite/-/wikis/api/api_reference

The connection of the visualisation tools with the *in situ* pipeline enables a direct view into a running simulation and, for instance, shortening the time to detect errors. It is the first step towards interactive simulator steering.

A middleware has been developed that can read data coming from running simulations through the *in situ* pipeline. This middleware has been integrated into the ViSimpl tool. The development has been carried out in close collaboration with the developers of the *in situ* pipeline to have a fast feedback loop for a quick integration as well as to deal with spikes in compartment-based simulations (see Figure 3).

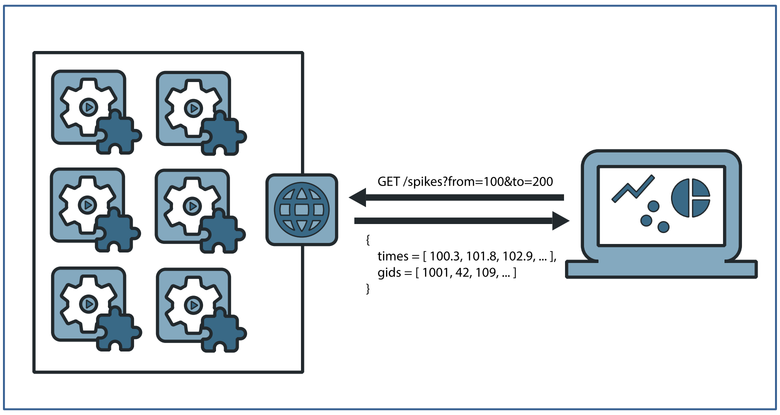

Figure 3: Restructured setup of the *in situ* pipeline

This figure shows the restructured setup of the in situ pipeline. A client connects to a single access point of the pipeline using standardised HTTP requests. The requested data is gathered from simulator-specific modules integrated into the simulation engine running on multiple compute nodes.

The in situ pipeline underwent a major rework in the second year of SGA2 to reduce the integration effort into visualisation tools. The rework was done in close collaboration with the users from URJC / UPM and included a 3-month research stay of Simon OEHRL at the URJC Madrid (as one action item identified in terms of C2698). In addition, Sebastian SPREIZER, developer of NEST Desktop (HBP Voucher #25) was included in the design decisions for the pipeline. All the necessary software is available inside Docker containers for easy deployment. Computing resources on the CSCS cluster have been granted for intensive testing of the pipeline.

6.1.3 Output 2: Deployment and extension of visualisation tools at CINECA

In collaboration with Michele MIGLIORE's group, two of the visualisation tools (ViSimpl and NeuroTeshMess, C2681) were deployed on GALILEO, a supercomputer at CINECA (see SGA2 Deliverable D7.6.1 (D47.1 D22) for details). The evaluation of the difficulties encountered, as well as the lessons learned during this experience, led this work towards the creation of Docker containers that will make it easy to deploy these tools (C2681) on any supercomputer, in EBRAINS and also outside of the HBP. The team of Colin MCMURTRIE at CSCS is currently testing these containers.

6.1.4 Output 3: Release of a prototype for remote visualisation and interaction

Table 14: KR7.4 Output 2 links

Component	Link to	URL
C2682	Software Repository	https://github.com/gmrvvis/ReMo
	Technical Documentation	https://github.com/gmrvvis/ReMo
	User Documentation	https://github.com/gmrvvis/ReMo

The developed remote rendering middleware allows running any application that has been built on the machines where the middleware is deployed (see SGA2 Deliverable D7.6.1 (D47.1 D22) for details). The entry point for the back end can be a web page or a Jupyter Notebook that allows users to create and manage sessions. The middleware is fully usable from the HBP Collaboratory and also allows to retrieve data from the Collaboratory and make it available in the machine where the back end is running. The connection with the archival database repository using swiftfs could, due to the Covid-19 crisis, not be finished during SGA2 and is therefore still work in progress.

6.1.5 Output 4: Visualisation software catalogue design and first version

Table 15: KR7.4 Output 4 links

Component	Link to	URL
C2707	Curation Repository	https://github.com/bweyers/HBPVisCatalogue/issues
	Technical Documentation	https://github.com/bweyers/HBPVisCatalogue/wiki
	User Documentation	https://github.com/bweyers/HBPVisCatalogue/wiki

As outlined in SGA2 Deliverable D7.6.1 (D47.1 D22), the identified shift from a standalone, newly developed semantic search engine towards adding searchable software metadata into the EBRAINS Knowledge Graph (KG) has been finalised. In collaboration with Andrew DAVISON (SP5) and the Data Curation Team, partner UT was able to develop a software metadata schema based on the previous investigation and developments that were conducted in the first half of SGA2. All details on the metadata schema can be found in the catalogue's documentation on GitHub and in SGA2 Deliverable D7.4.1 (D45.1 D84). With the help of UT, the schema has been implemented by the KG development team (led by Oliver SCHMID, SP5) in the KG as well as in the KG editor. The release of the first version of the software metadata card took place together with the release of the new front end at the start of the HBP Summit in Athens.

In close collaboration (weekly data curation videoconferences, in which UT started to participate) and with support by the Data Curation Team, UT established curation processes for the step-wise research and integration of software metadata into the Knowledge Graph from various sources. All processes are documented in the GitHub repository and described in detail in SGA2 Deliverable D7.4.1 (D45.1 D84). The curation processes are transparently documented in the GitHub repository's issue tracker (https://github.com/bweyers/HBPVisCatalogue/issues, 15 fully integrated into Knowledge Graph & released, 40 data sets in curation, status of 2020-02-11). UT has started to integrate software metadata from the Collaboratory software catalogue, and of the HPAC Guidebook following suggestions from the Project. However, the catalogue will not only consider software developed within the HBP but also list software relevant for neuroscience and/or provided on EBRAINS resources.

For broader coordination and face-to-face discussions, UT organised together with Andrew DAVISON a meeting in Heidelberg during the CodeJam#10 (Dissemination 149 in PLUS). During this meeting, a subset of the PLUS, Collaboratory, KG, and Data Curation Team members were present to discuss overlaps and potential integration. In addition, UT participated in the data curation helpdesk during the HBP Summit to advertise the software metadata curation to the Project Consortium.

6.2 Validation and Impact

6.2.1 Actual and Potential Use of Output(s)

The visualisation tools (C2681) are being used by the groups of Michele MIGLIORE (CNR), Idan SEGEV (HUJI), Egidio D'ANGELO (UNIPV), Eduardo ROS (UGR), the SimLab Neuroscience (JUELICH-JSC), and Simo VANNI (University of Helsinki and Helsinki University Hospital). The tools are at TRL7 and could be exploited by HBP users that are in need of visual interactive analysis and data exploration. The techniques developed are domain-independent, so that they can be used for other visualisation problems outside of the HBP. The Docker containers developed facilitate the deployment of the tools for different types of users.

The remote visualisation middleware (C2682) was already released (MS7.3.13) and is currently at TRL5. It can be used for all use cases where visualisation is needed, not only in the HBP and the neuroscience domain.

The *in-situ* pipeline (C2680) is currently used in the visualisation tools (2681) and is planned to be integrated into NEST Desktop. The pipeline will be also adapted to the needs of the co-simulation workflow in SGA3.

Software metadata (beyond visualisation software as planned for C2707) is now searchable in the KG and can be used as entities in terms of advanced semantic descriptions, e.g. in terms of provenance tracking, in relation with curated data sets, as well as for making software searchable from outside of HBP as well (KG search interface is openly accessible). The original goals of making visualisation software tools available and semantically searchable has been achieved due to the use of the established software stack provided by the EBRAINS KG, and extended towards general software. Additionally, the metadata gathering process is supported by a transparent metadata curation process.

6.2.2 Publications

No publications yet.

7. Key Result KR7.5 High Performance Analytics and Computing Platform v3

We provide the base infrastructure and infrastructure services for the platforms of the HBP. Compute and storage resources are provided via the ICEI resource allocation mechanisms and there are a number of projects running on the infrastructure, using primarily resources at ETHZ-CSCS and also the pilot system JURON at JUELICH-JSC.

In the last few months, progress was made in a number of areas. This included the integration of the CEA into the HPAC Platform, the development of the container runtime engine Sarus tailored to HPC systems, continued support for the Neurorobotics Platform, and also increased collaboration with SP5.

In addition, the weekly HBP Joint Infrastructure Coordination meetings are proving to be ever more useful and productive as a way to implement a use-case driven approach. These meetings are well-attended across various Platforms, with a GitLab setup and agile approach. This is important to ensure close contacts and for finding solutions to issues identified that concern multiple Platforms. To date, over 100 issues have been closed, highlighting the effectiveness and usefulness of this approach.

7.1 Outputs

7.1.1 *Overview of Outputs*

7.1.1.1 List of Outputs contributing to this KR

- Output 1: Integration and Operation of Low-level Infrastructure (C263, C329, C407, C409, C410, C416, C417, C571, C1092, C1094, C1096, C1097, C1112, C1129, C2828, C2831, C2832, C2834, C2837-C2843, C2849, C2855, C2860, C2861, C2870, C2875, see Table 15)
- Output 2: Policy management (C2871, C2872, C2873)
- Output 3: Infrastructure services (C263, C329, see Table 16)
- Output 4: Platform developer services (C329)
- Output 5: Platform integration services (C263, see Table 17)
- Output 6: Performance optimisation services (C2709)
- Output 7: User support and documentation (C337, C1114, see Table 18)
- Output 8: Education and training (C2867)

7.1.1.2 How Outputs relate to each other and the Key Result

The Outputs of KR7.5 High Performance and Computing Platform v3 cover work of Tasks in WP7.4 *User support and community building* as well as work of the whole WP7.5 *Integration and operation*. KR7.5 is associated with SO7.5 *Establishment of advanced user support services and community building* and SO 7.6 *Integration of CEA as fifth HPC and data centre into the HPAC Platform*. The HPAC Platform provides neuroscientists with the high performance computing, storage and data processing capabilities they need to run simulations of sophisticated, detailed brain models and to analyse large, complex data sets.

- Output 1 Integration and Operation of Low-level Infrastructure contributes to KR7.5 by maintaining and coordinating operations of the federated HPAC infrastructure including network, compute, storage and monitoring services.
- Output 2 Policy Management contributes to KR7.5 and builds on Output 1 Integration and Operation of Low level Infrastructure by updating and enforcing security policies related to the federated infrastructure and also defining resource allocation policies for the federated infrastructure.
- Output 3 Infrastructure services contributes to KR7.5 by engaging with the HBP HLST to provide support to identify and resolve issues that concern multiple Platforms using the federated infrastructure.
- Output 4 Platform developer services contributes to KR7.5 and builds on Outputs 1,2 and 3 by supporting virtualised environments for end users on the federated infrastructure.
- Output 5 Platform integration services contributes to KR7.5 by maintaining and operating the middleware, including workflow management, allowing end users to access the federated infrastructure from the HBP Collaboratory.
- Output 6 Performance optimisation services contributes to KR7.5 by supporting users to optimise the performance of their applications and migrating codes to new architectures.
- Output 7 User support and documentation contributes to KR7.5 by providing general user support for the HPAC platform.
- Output 8 Education and training contributes to KR7.5 by providing training opportunities related to the HPAC platform.

7.1.2 Output 1: Integration and Operation of Low-level Infrastructure

Table 16: KR7.5 Output 1 links

Component	Link to	URL
C263	Technical Documentation	https://www.unicore.eu/
	User Documentation	https://www.unicore.eu/
C416	Technical Documentation	https://www.unicore.eu/docstore/uftp-2.2.0/uftp-manual.html
	User Documentation	https://collab.humanbrainproject.eu/#/collab/3656/nav/275433

This table only lists relevant components with releases and the corresponding technical or user documentation. Not listed here are components with no separate releases or components that are continuous activities by nature and thus have no releases.

The operation of low-level infrastructure at the partner sites (BSC, CEA, CINECA, ETHZ-CSCS, JUELICH-JSC, KIT) comprising the compute-, data- and network infrastructure services, has been coordinated. Day-to-day operational activities, such as configuration and software update management, have been performed to ensure a continuously high availability of the low-level services. At ETHZ-CSCS, this encompassed also the support for the approved ICEI projects and associated users. Development work is on-going to provide the Platforms, such as the HPAC Platform, with an improved visibility about scheduled and unscheduled maintenances on the infrastructure.

The integration of the CEA infrastructure has been successfully completed. To this end, a method has been implemented to allow users to perform a local authentication step at CEA, which is required to meet legal obligations, and which is also used to generate and locally store a Kerberos ticket. This ticket is automatically and securely retrieved by the UNICORE daemon, thus enabling users of the HPAC platform to seamlessly use the CEA resources. SGA2 Deliverable D7.5.1 (D46.1 D85) documented an example workflow.

7.1.3 Output 2: Policy Management

Table 17: KR7.5 Output 2 links

Component	Link to	URL
C2871 C872 C2873	Public Milestone Report	https://collab.humanbrainproject.eu/#/collab/264/nav/1975?state=uuid%3D59da41e3-43a0-478c-9873-f095cb8314af

A review of security policies is underway, to understand the relationship between site-local security policies and the overall security of the federated infrastructure. This is intimately tied to the ICEI project. The output of the Data Governance Working Group (DGWG), namely the Data Management Plan (DMP), is also currently under review.

7.1.4 Output 3: Infrastructure Services

Table 18: KR7.5 Output 3 links

Component	Link to	URL
C263	Technical Documentation	https://www.unicore.eu/
	User Documentation	https://www.unicore.eu/

Work with HBP HLST to support other SPs was continued. Weekly HBP Joint Infrastructure Coordination meetings were used as a means for identifying and resolving issues that concern multiple Platforms. Furthermore, work on the Data Transfer Service was engaged, resulting in discussions between CINECA, JUELICH-JSC and ETHZ-CSCS. Finally, investigation of containerisation of visualisation tools has recently started.

Output 4: Platform developer services

The “Lightweight Virtualisation Report”[6], provides an overview of ‘lightweight virtualisation’, which demands less resources from the underlying host system, has better responsiveness and can leverage native or close-to-native performance from the actual host system. In particular, information about the production deployment of Sarus as a replacement for Shifter at CSCS is detailed. Sarus leverages the OCI standards to offer modularity and extensibility, and offers system administrators the possibility to configure OCI hooks: these are standalone programs which can customise a container, thus extending the functionality of an OCI runtime by working as plugins. Sarus has been deployed in production at CSCS since November 2019.

On the Neurorobotics Platform (NRP) side, ETHZ-CSCS is continuing to support developments in SP10, with a view to integrating Piz Daint into the NRP workflow. Significant progress has been made in this area, and we expect to be able to run scientific use cases from the University of Granada and the University of Pavia soon to be able to test the execution of NRP and distributed NEST.

7.1.5 Output 5: Platform integration services

Table 19: KR7.5 Output 5 links

Component	Link to	URL
	Technical Documentation	https://www.unicore.eu/
	User Documentation	https://www.unicore.eu/

A new major release (UNICORE 8.0) of the UNICORE software suite has been finished, which focuses on simpler deployment, administration and maintenance of the software, as well as implementing a number of features requested by HBP Platform users.

The authentication and authorisation infrastructure that allows users to use HPC resources from the Collaboratory via UNICORE has been updated to support the new HBP identity provider (https://iam.humanbrainproject.eu/auth/realms/hbp/account).

The usage of UNICORE by other groups in the HBP was continuously supported. For example, UNICORE is used heavily by the Brain Simulation Platform to access HPC resources at JUELICH-JSC and ETHZ-CSCS from the HBP Collaboratory.

7.1.6 Output 6: Performance Optimisation

Table 20: KR7.5 Output 6 links

Component	Link to	URL
C2709	Technical Documentation	https://collab.humanbrainproject.eu/#/collab/264/nav/1975?state=uuid%3Dba213669-48b4-430a-b0f6-5fc44a62d6a8
	User Documentation	https://collab.humanbrainproject.eu/#/collab/264/nav/1975?state=uuid%3Dba213669-48b4-430a-b0f6-5fc44a62d6a8

The work on performance analysis and optimisation of the second year of SGA2 has focused on evaluating the new version of CoreNEURON provided by the developers. The latest version uses NMODL, a new source-to-source compiler, to process modelling files. With this framework, the program generates optimised source files from original modelling files, which one may expect to

[6] https://collab.humanbrainproject.eu/#/collab/264/nav/1975?state=uuid%3D8df2b4a8-4f47-4f93-91af-4aac997ab0b4

result in overall performance improvement. A performance comparison of the new version using NMODL against the old one has been performed and reported to CoreNeuron developers.

Also, there was a request from CoreNEURON developers to study the performance of the new code in an Arm-based cluster. To fulfil this request, CoreNEURON was deployed in an Arm-based platform, evaluated there, and compared the performance achieved with an Intel-based platform.

7.1.7 Output 7: User support and documentation

Table 21: KR7.5 Output 7 links

Component	Link to	URL
C337	Technical Documentation	https://hbp-hpc-platform.fz-juelich.de/
	User Documentation	https://hbp-hpc-platform.fz-juelich.de/
C1114	Technical Documentation	https://collab.humanbrainproject.eu/#/collab/264/nav/329319
	User Documentation	https://collab.humanbrainproject.eu/#/collab/264/nav/329319

The HPAC user support Task has been focused during the second year of SGA2 on providing the first-level support for the HPAC infrastructure and improving the interaction with the central HBP user support ticket system (https://support.humanbrainproject.eu/). Furthermore, the interaction with the teams providing HLST support and 2nd level support within HBP has been improved, with a better understanding of the different responsibilities and with the creation of dedicated HLST queues per site. The current workflow is designed such that the incoming tickets can be assigned to 1st level support, 2nd level support teams or HLST depending on the topic, site and level of request. The documentation of the HPAC Platform is collected in the Guidebook, which has been updated on a regular basis during the full project duration.

7.1.8 Output 8: Education and training

The work on education and training focused on organising the 2nd HPAC Platform Training Event (E152, https://plus.humanbrainproject.eu/events/152/). During the planning phase, we realised that other related events were scheduled for a similar period of time (end of 2019), and thus we decided to join efforts to create a joint event together including also a public EBRAINS Event, CodeJam #10 and an HLST Workshop (HBP High-Level Support Team).

The agenda of the 2nd HPAC Platform Training was based on the feedback, successes and lessons learned of the 1st HPAC Platform Training Event and was aligned with the programme of the other events taking place in parallel, so that participants of one event could also attend sessions of the other events that were relevant for them.

The joint event took place from 26-28 November 2019 in Heidelberg, Germany. In total, there were 26 registrants participating the HPAC Platform Training, some of them external to HBP.

The topics covered during the training included an introduction to the resources, tools and services provided by the HPAC Platform and by the Fenix infrastructure, the different options to get access to HPC and data resources, as well as how to use Fenix resources and services in workflows. It also included hands-on sessions on how to transfer between HPC sites and how to access resources from Jupyter notebooks in the HBP Collaboratory. The participants also got an overview of the visualisation tools available in the HBP. The simulator NEST, for models with point-neurons, was introduced, including more in-depth sessions introducing NEST Desktop, NESTML and an outlook to NEST 3. Another session dealt with the coupling of the simulators NEST and TVB.

7.2 Validation and Impact

7.2.1 Actual and Potential Use of Output(s)

The HPAC Platform provides access to the base infrastructure and infrastructure services for the Platforms of the HBP. Compute and storage resources are provided via the ICEI resource allocation mechanisms. An increase in the number of projects allocated to HBP members from 8 in M12, to the current total of 23, excluding those under review, gives an indication of the validation and impact. These projects are reporting good results, and a number have already re-applied for resources as they clearly appreciate the benefit that they gain. These projects are from across different areas of HBP, clearly demonstrating the benefit to the HBP as a whole. In addition, the on-boarding of Fenix resources from other providers apart from ETHZ-CSCS means that there will be an increase in the capacity of resources that are offered, benefitting even more users.

7.2.2 Publications

No publications yet.

8. Conclusion and Outlook

The HPAC Platform has made important contributions towards the establishment of a sustainable research infrastructure for the neuroscience community. It works very closely with the other HBP Subprojects and CDPs to support their science and use cases with HPAC tools and services, as well as access to Fenix resources.

An architecture of what so far was called the "HPAC Platform" has been specified, based on the generic concept that basic e-infrastructure services are provided by Fenix through the ICEI project, and neuroscience-specific Platform services, which are running on top of these e-infrastructure services, are designed, deployed and operated as part of SGA2 (and, in the future, SGA3). Details of these services continue to be worked out, taking the results of use case analyses into account. To validate architecture and implementation details, a set of validation tests have been formulated and are being implemented. To further exploit the generated knowledge, efforts have been made to transfer this knowledge to HPC and Cloud solutions providers, to improve their understanding of the needs of the HBP and the future EBRAINS Research Infrastructure.

Several important technologies for data federation and data-intensive computing were developed in the last year. The Fenix data transfer service has been enabled and PCOCC has been advanced, which allows users to host their own cluster of virtual machines on HPC compute nodes alongside regular HPC jobs. In addition to that, work on key technologies for data access and retrieval, as well as on a SLURM plug-in for the co-allocation of compute and data resources, was implemented.

The NEST simulation technology has advanced significantly during SGA2, including the addition of numerous new features such as complex learning rules and high-profile neuron models. The development of a powerful new user interface for network construction in combination with significant simulation kernel refactoring for long-term stability has progressed far, with a release of NEST 3.0 in the near future. In parallel, NEST documentation has been improved strongly, and the new NEST Desktop user interface provides much easier access to network simulations, especially for education.

In SGA3, the support for third-factor learning rules will be extended, the NEST simulation kernel will be revised to remove bottlenecks of performance characterised in SGA2 and to make use of the vectorisation potential of upcoming conventional computer architectures. Documentation and training material will be improved and extended following the detailed concept developed in SGA2. In addition, after a long phase of rapid development to meet the requirements of HBP on functionality and performance, the NEST code will be systematically refactored to ensure long-term sustainability. In SGA3, the NESTML code generation for synapses will be extended to include third-factor plasticity rules, enabled by SGA2 accomplishments in the NEST Simulator. Maintainability of the NEST Simulator codebase will be improved by tighter integration with NESTML.

The Arbor developers achieved all of their objectives in a timely manner. The main lesson learnt during the development was the complexity of implementing a robust interface for users to describe arbitrarily complex models and workflows, which was underestimated during initial planning. The Arbor simulation engine is ready for integration into the EBRAINS platform. To support Arbor, models will need to be ported from NEURON to simulator-agnostic formats, which in turn will make the platform more accessible. By virtue of running on CPU and GPU HPC systems, Arbor will also increase flexibility for both users, who will be able to run on more systems, and HPC centres to provide Arbor on more of their systems.

The development of TVB-HPC will allow efficient fitting of TVB simulations, leveraging the computing power of the HBP infrastructure, which can be used in the future for innovative clinical applications with a potential impact in society.

In SGA3 we will further strengthen the strong NEST community engagement developed over the duration of the HBP in close collaboration with the HLST and Community Building Tasks in SGA3 WP4, while documentation will be developed further as part of our efforts in SGA3 WP5 in collaboration with the HLST. As EBRAINS becomes fully available as a research Platform for computational neuroscience, we will work with users around Europe and beyond to enable cutting-edge neuroscience characterised by collaboration and FAIR model sharing and re-use in an expanding EBRAINS and NEST community.

During SGA2, simulator and visualisation development in SP7 have been more tightly connected than in previous phases, which enabled a close collaboration between both topics. The *in-situ* pipeline has been continuously extended to support NEST, Arbor and TVB. The connection of the visualisation tools with the *in-situ* pipeline enables a direct view into a running simulation and, for instance, shortening the time to detect errors. It is the first step towards interactive simulator steering. In terms of the visualisation catalogue, some adjustments have been made to the original Work Plan to broaden the scope and use of the gathered metadata towards general software. However, the major goals have been addressed and reached in close collaboration with the users.

Also, the HPAC Platform itself has seen major changes in SGA2. Most notably, the migration towards the new Fenix Infrastructure and the integration of the TGCC of CEA as fifth supercomputing centre into the HPAC Platform and Fenix. Overall, the vast majority of targets were achieved, albeit with some delay compared to the initial work plan. One reason for this was that the technical integration of the CEA was significantly more complicated than anticipated in the planning of SGA2. This delay, however, did not have an impact on other parts of the HBP, and TGCC as a new infrastructure component for the HPAC Platform is now available. Another major Output that was completed in the last months was the development of Sarus as replacement for Shifter at CSCS and as a means for deploying containers at scale.

One of the main lessons learned regarding software development and deployment environments is the importance of a close user-developer interaction as well as a strong communication within the HBP. We faced the challenge of limited uptake of development and deployment services, as various teams are making use of their own tools. However, more dissemination and outreach are needed and planned to increase the uptake of such services.

To increase the number of users from inside and outside of the HBP, the HPAC Platform has increased its dissemination and training activities, often as joint efforts together with the ICEI project. Two HPAC Platform Trainings were organised for interested users during SGA2 that took place end of 2018 and 2019. In addition, an HBP Education Workshop “High Performance Computing for Neuroscience” was organised in July 2019.

In close collaboration, HPAC Platform partners and Fenix will work together beyond SGA2 towards consolidating and operating the basis for the HBP Research Infrastructure EBRAINS in SGA3, and providing simulation engines and visualisation software optimised for current and future HPC architectures.