SP7 High-Performance Analytics and Computing Platform - Results for SGA2 Year 1 (D7.6.1 - SGA2)

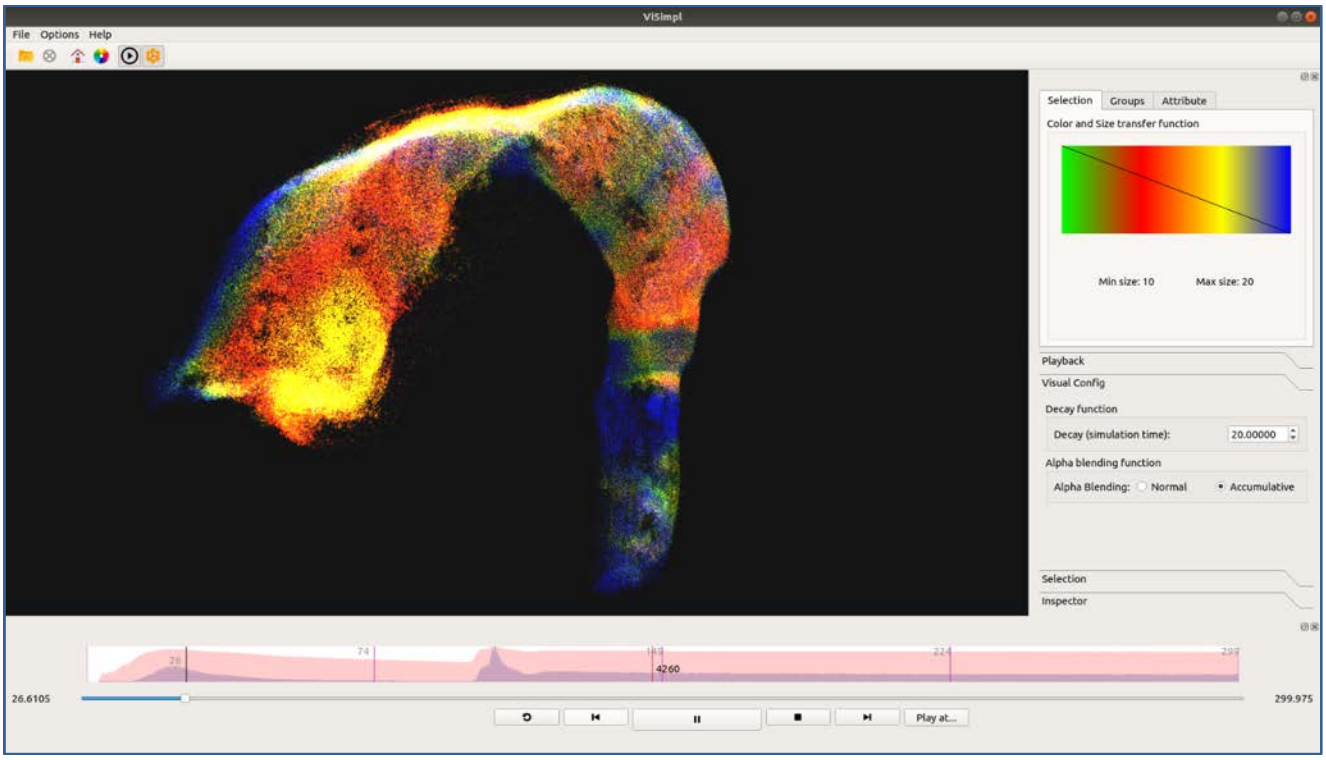

Figure 1: Visualisation of a hippocampus model using ViSimpl

This model of the hippocampus has been simulated by Michele Migliore (on behalf of the Hippocampus Team of HBP Brain Simulation Platform (SP6; EPFL, CNR, IEM HAS, UCL) using HPC resources at CINECA. The visualisation with ViSimpl (see Section 6 for more details) has been rendered on the GALILEO HPC system at CINECA.

Project Number:	785907	Project Title:	Human Brain Project SGA2

Document Title:	SP7 High-Performance Analytics and Computing Platform – Results for SGA2 Year 1
Document Filename:	D7.6.1 (D47.1 D22) SGA2 M12 ACCEPTED 200731.docx
Deliverable Number:	SGA2 D7.6.1 (47.1, D22)
Deliverable Type:	Report
Work Package(s):	WP7.1, WP7.2, WP7.3, WP7.4, WP7.5, WP7.6
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA2 M12 / 31 Mar 2019, Request for revision: 22 Jul 2019
Actual Delivery Date:	Submitted: SGA2 M12 / 26 Mar 2019, Resubmitted: SGA2 M21 / 10 Dec 2019; Accepted: 31 Jul 2020

Authors:	Thomas LIPPERT, JUELICH (P20), SP Leader Anna LÜHRS, JUELICH (P20), SP Manager
Compiling Editors:	Anna LÜHRS, JUELICH (P20), SP Manager
Contributors:	Javier BARTOLOMÉ, BSC (P5) Lucas BENEDICIC, ETHZ (P18) Matthias BOLTEN, BUW (P7) Juan Pedro BRITO, UPM (P68) Benjamin CUMMING, ETHZ (P18) Sandra DIAZ, JUELICH (P20) Markus DIESMANN, JUELICH (P20) Anne ELFGEN, JUELICH (P20) Giuseppe FIAMENI, CINECA (P13) Marta GARCIA GASULLA, BSC (P5) Thorsten HATER, JUELICH (P20) Wouter KLIJN, JUELICH (P20) Dorian KRAUSE, JUELICH (P20) Torsten KUHLEN, RWTH (P46) Stefanie LACKNER, TUDA (P85) Jacques-Charles LAFOUCRIERE, CEA (P11) Charl LINSSEN, JUELICH (P20) Susana MATA, URJC (P69) Colin MCMURTRIE, ETHZ (P18) Jessica MITCHELL, JUELICH (P20) Abigail MORRISON, JUELICH (P20) Luis PASTOR, URJC (P69) Alexander PEYSER, JUELICH (P20) Dirk PLEITER, JUELICH (P20)

	Hans Ekkehard PLESSER, NMBU (P44) Maria RIBERA SANCHO, BSC (P5) Sebastian RINKE, TUDA (P85) Oscar David ROBLES, URJC (P69) Marco RORRO, CINECA (P13) Bernd SCHULLER, JUELICH (P20) Dennis TERHORST, JUELICH (P20) Debora TESTI, CINECA (P13) Pablo TOHARIA, UPM (P68) David VICENTE, BSC (P5) Benjamin WEYERS, UT (P130) (until Dec. 2018 RWTH, P46) Gilles WIBER, CEA (P11) Gabriel WITTUM, UFRA (P35) Felix WOLF, TUDA (P85) Stuart YATES, ETHZ (P18)
SciTechCoord Review:	Yannick MOREL, TUM (P56); Revision: Mehdi SNENE, EPFL (P1)
Editorial Review:	Guy WILLIS, EPFL (P1); Revision: Guy WILLIS, EPFL (P1)

Description in GA:	High Performance Analytics and Computing Platform v2.5. The Platform, accessible to end users through the HBP Collaboratory, will provide the HBP and the neuroscience community with access to federated data and scalable compute resources as well as visualisation and simulation technology.
Abstract:	This Deliverable summarises the Key Results of the High-Performance Analytics and Computing (HPAC) Platform in the first year of SGA2. These comprise, among others, a first 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 Contents 4
Table of Tables 5
Table of Figures 5
1. Overview 7
2. Introduction 7
3. Key Result KR7.1 Architecture specification for the HPAC Platform and Fenix 8
3.1 Outputs 9
3.1.1 Overview of Outputs 9
3.1.2 Collection of science and use case documents 9
3.1.3 "Approach for platform validation for SGA2 defined" Report 9
3.1.4 "Knowledge transfer process for SGA2 defined" Report 10
3.2 Validation and Impact 10
3.2.1 Actual Use of Output(s) / Exploitation 10
3.2.2 Potential Use of Output(s) 10
3.2.3 Publications 10
3.2.4 Measures to Increase Impact of Output(s): Dissemination 10
4. Key Result KR7.2 Data federation and data-intensive computing technology 11
4.1 Outputs 11
4.1.1 Overview of Outputs 11
4.1.2 HBP/Fenix data transfer service 11
4.1.3 Private Cloud On a Compute Cluster (PCOCC) 13
4.2 Validation and Impact 14
4.2.1 Actual Use of Output(s) / Exploitation 14
4.2.2 Potential Use of Output(s) 14
4.2.3 Publications 14
4.2.4 Measures to Increase Impact of Output(s): Dissemination 14
5. Key Result KR7.3 Exascale-ready simulation technology: NEST, Arbor and TVB 15
5.1 Outputs 15
5.1.1 Overview of Outputs 15
5.1.2 NEST 16
5.1.3 Arbor 16
5.1.4 TVB 17
5.1.5 Multi-simulator multi-scale interaction 17
5.2 Validation and Impact 17
5.2.1 Actual Use of Output(s) / Exploitation 18
5.2.2 Potential Use of Output(s) 18
5.2.3 Publications 19
5.2.4 Measures to Increase Impact of Output(s): Dissemination 19
6. Key Result KR7.4 Interactive visual data analytics and in-situ visualisation developed and deployed on Fenix and HPAC Platform infrastructure 20
6.1 Outputs 20
6.1.1 Overview of Outputs 20
6.1.2 In Situ Pipeline 20
6.1.3 Deployment and extension of visualisation tools at CINECA 21
6.1.4 Release of a prototype for remote visualisation and interaction 22
6.1.5 Hands-on workshop and roadmap definition 23
6.1.6 Visualisation Software Catalogue 24
6.2 Validation and Impact 25
6.2.1 Actual Use of Output(s) / Exploitation 25
6.2.2 Potential Use of Output(s) 25
6.2.3 Publications 26

6.2.4 Measures to Increase Impact of Output(s): Dissemination 26
7. Key Result KR7.5 High Performance Analytics and Computing Platform v3 26
7.1 Outputs 27
7.1.1 Overview of Outputs 27
7.1.2 Integration and Operation of Low-level Infrastructure 27
7.1.3 Policy Management 28
7.1.4 Infrastructure Services 28
7.1.5 Platform developer services 29
7.1.6 Platform Integration Services 30
7.1.7 Performance optimisation 30
7.1.8 User support and documentation 31
7.1.9 Education and training 31
7.2 Validation and Impact 31
7.2.1 Actual Use of Output(s) / Exploitation 31
7.2.2 Potential Use of Output(s) 32
7.2.3 Publications 32
7.2.4 Measures to Increase Impact of Output(s): Dissemination 32
8. Conclusion and Outlook 32
Annex A: Additional information about releases and major updates 34

Table of Tables

Table 1: Overview of outputs related to Key Result KR7.1 9
Table 2: Overview of outputs related to Key Result KR7.2 11
Table 3: Overview of outputs related to Key Result KR7.3 15
Table 4: Overview of outputs related to Key Result KR7.4 20
Table 5: Overview of outputs related to Key Result KR7.5 27
Table 6: Overview of releases and major updates related to Key Result KR7.1 34
Table 7: Overview of releases and major updates related to Key Result KR7.2 34
Table 8: Overview of releases and major updates related to Key Result KR7.3 35
Table 9: Overview of releases and major updates related to Key Result KR7.4 37
Table 10: Overview of releases and major updates related to Key Result KR7.5 38

Table of Figures

Figure 1: Visualisation of a hippocampus model using ViSimpl 1
Figure 2: General view of the data transfer services in the HPAC Platform and Fenix 12
Figure 3: PCOCC general architecture 13
Figure 4: Simulation of the Hippocampus in ViSimpl with selection of different groups of neurons 21
Figure 5: Morphologies of a region from the hippocampus viewed in NeuroTessMesh 22
Figure 6: Scheme of components of the remote rendering framework 23
Figure 7: Circular layout of a network created using ConGen 24
Figure 8: CSCS Jenkins CI service in use for Arbor development 30

History of Changes made to this Deliverable (post Submission)

Date	Change Requested / Change Made / Other Action
26 Mar 2019	Deliverable submitted to EC
22 Jul 2019	Resubmission with specified changes requested in Review Report Main changes requested: • Change 1: The roadmap for the integration of in-situ visualisation with data analytics tools (which tools will be coupled with in situ visualisation, which tools are planned to be newly developed) [...] should be included. • Change 2: [...] a justification of selection of the visualisation tools should be included. • Change 3: A clear description and strategy for the Visualisation Software catalogue should be added
8 Nov 2019	Revised draft sent by SP7 to PCO. Main changes made, with indication where each change was made: • Change 1: see Sections 6.1.4 and 6.1.5 • Change 2 see Section 6.2.1 • Change 3: see Section 6.1.6
10 Dec 2019	Revised version resubmitted to EC by PCO via SyGMa

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.

A significant change in the last year was the start of the Interactive Computing E-Infrastructure (ICEI) project, funded through a separate HBP Grant Agreement, which builds the Fenix[1] infrastructure. Fenix and the HPAC Platform are tightly connected, as HPAC services run on Fenix infrastructure and the HPAC Platform also develops some of the Fenix services. Both projects are implemented and operated by the same five supercomputing centres. The HPAC Platform has started to migrate its services to use the Fenix infrastructure instead of legacy resources at the involved centres. HPAC also develops technologies for Fenix, e.g. the data transfer service, and it investigates technologies to support interactive data analysis and visualisation. Another important, ongoing effort is the integration of CEA's TGCC[2] in France as the fifth High-Performance Computing (HPC) centre in the HPAC Platform. With this new addition, HPAC and Fenix can now offer resources and services in five European countries; all five HPAC sites are also PRACE hosting sites.

The HPAC Platform is also 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 were achieved in the past year. Also, progress has been made in the coupling of simulators and a benchmark and validation suite for simulation engines has been developed. In addition, interactive visual data analytics and *in-situ* visualisation tools and techniques are under development, with a focus on the visualisation of simulation data. Work on a visualisation software catalogue has started, which will help scientists to identify which tools are most suitable for their purposes.

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 the last year was the start of the Interactive Computing E-Infrastructure (ICEI) project, funded through a separate HBP Grant Agreement, which builds the Fenix infrastructure. ICEI and the HPAC Platform are tightly connected, as they are implemented and operated by the same five supercomputing centres. France's CEA joined the HPAC Platform at the start of SGA2 and the technical integration of its *Très Grand Centre de calcul* (TGCC) HPC facility is still ongoing. The HPAC Platform has started to migrate its services to run on the Fenix infrastructure, instead of using legacy resources at the involved centres (see Key Result 7.5, Section 7). Some HBP research groups have already started using Fenix 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 development of an aligned architecture specification for the HPAC Platform and Fenix, based on a detailed analysis of use cases and derived requirements, was important for ensuring the successful migration to the Fenix infrastructure (see Key Result 7.1, Section 3). In addition, an approach for the validation of the HPAC Platform, based on the use cases and a knowledge transfer process for facilitating the sharing of knowledge and results with external stakeholders, has been defined to maximise the impact of our results beyond the HBP.

[1] https://fenix-ri.eu/
[2] Très Grand Centre de calcul (TGCC), http://www-hpc.cea.fr/en/complexe/tgcc.htm

The HPAC Platform also develops technologies for Fenix (see Key Result 7.2, Section 4), e.g. the data transfer service and PCOCC, a technology for hosting clusters of virtual machines on compute nodes alongside regular High-Performance Computing (HPC) jobs. Investigation of technologies to support interactive data analysis and visualisation has started.

Another area where significant efforts were invested was the improvement of simulation software and work to ensure their readiness for current and future HPC systems (see Key Result 7.3, Section 5). NEST version 2.16.0 has been released, which comprises a number of novel features and improvements that prepare this simulator for future (pre-) Exascale HPC systems. NESTML, the high-level modelling language developed for NEST, now supports additional features. The Arbor simulation software had its first major release and is being continuously updated and improved. In addition, a benchmark and validation suite for simulation engines has been released. In the context of CDP8, SP7 has started contributing to TVB and has already provided an initial design for the Domain-Specific Language for neural mass models. Moreover, work has started to prepare multi-simulator, multi-scale simulations.

Other important work undertaken within the HPAC Platform includes interactive visual data analytics and *in-situ* visualisation tools and techniques (see Key Result 7.4, Section 6). The focus in SGA2 is on the visualisation of simulation data. In the first year, efforts have been concentrated on an *in-situ* visualisation pipeline, as well as the release of a prototype for remote visualisation and interaction. Another important activity was the deployment and extension of visualisation tools at CINECA for a group from SP6/CDP4. Work on a visualisation software catalogue has started, which will help scientists to identify which tools, either those available within the HBP or open source, are most suitable for their purposes.

To ensure that users can get the most out of the HPAC Platform, the HPAC Guidebook[3], which collects all relevant documentation, has been continuously updated. In complement, dissemination, education and training events have been organised, often as joint efforts between Fenix and HPAC. The HPAC team has supported scientists in their work, both through the central support team, the HLST and at a more advanced level, e.g. by the performance optimisation service.

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

The main objectives of WP7.1 “Architecture specification and validation” are to:

1) Provide a specification of the architecture of the HPAC Platform and related aspects of the ICEI infrastructure;
2) Collect and document use cases for co-design and validation of the HPAC Platform;
3) Execute a validation of the HPAC Platform;
4) Enable knowledge transfer to relevant HPC and High Performance Data Analysis (HPDA) solution providers.

The work concerning the architecture specification (objective #1) is ongoing, but publication of a first version became significantly delayed and Deliverable D7.1.1 “HPAC Platform Architecture Specification v1” has not been submitted yet. Internally, all services that are part of HPAC Platform have been defined, but integration into a coherent architecture compliant with the ICEI infrastructure services portfolio has not been finished yet. The delays were caused by late submissions and resubmissions of ICEI project Deliverables that D7.1.1 depends on. The delayed availability of the architecture specification has had no significant impact on other parts of SP7 or HBP as a whole, because all involved partners have been kept informed about the current planning status. Preliminary solutions have been found with and for all users concerned.

[3] https://hbp-hpc-platform.fz-juelich.de/

In a common effort with the HBP ICEI project, a portfolio of science and use cases has been collected. The goal was to identify all cases for potential use of the HPAC Platform and the underlying ICEI infrastructure services, i.e. the future EBRAINS Research Infrastructure[4].

To prepare for validation of the HPAC Platform services, a strategy has been defined and documented in an internal report. The same was done in the context of knowledge transfer. A characteristic that both strategies have in common is that application benchmarks play a central role. In a common effort with the ICEI project, a collection of benchmarks has been implemented and tested on different systems. The benchmarks are based on applications frequently used in the HBP, e.g. Arbor, NEST, NEURON and Elephant.

3.1 Outputs

3.1.1 Overview of Outputs

Table 1: Overview of outputs related to Key Result KR7.1

Output	Component number(s)	Component name(s)	Additional information
Collection of science and use case documents	C2671	SP7 use case and requirements documentation	
Report "Approach for platform validation for SGA2 defined"	C2672	HPAC Platform validation (SGA2)	Report provided to document that milestone MS7.1.5 has been reached
Report "Knowledge transfer process for SGA2 defined"	C2674	SP7 co-design knowledge management (SGA2)	Report provided to document that milestone MS7.1.6 has been reached

3.1.2 Collection of science and use case documents

In an effort to better understand the current and future requirements of potential users of the EBRAINS Research Infrastructure, a portfolio of science cases and use cases has been defined. For each of these, a living document has been produced. While the initial focus was on collecting infrastructure services requirements for the ICEI project, WP7.1 extended these efforts to also collect requirements for HPAC Platform services. The documentation followed a template for systematic collection of infrastructure services requirements following an approach which we called "annotated use case diagrams". This approach allows the overall workflow to be captured by mapping it onto different processing stations; along with data repositories, these are represented as nodes of the diagram. By analysing the edges in the diagrams, an understanding of the required data transport capabilities within the infrastructure can be obtained.

In an ongoing next step, the science and use case documents are being systematically augmented to capture the needs with respect to HPAC Platform services requirements.

3.1.3 "Approach for platform validation for SGA2 defined" Report

The aim of this report was to define how the validation of the HPAC Platform will be carried out, as soon as all services are ready to support the execution of the identified use cases. The validation effort will be organised through an iterative approach, by adding more and more tests and

[4] This is not to be confused with the collection of the strategic use cases maintained by the project to steer the future work.

benchmarks to a validation test portfolio. A template has been developed for defining these tests and benchmarks.

3.1.4 "Knowledge transfer process for SGA2 defined" Report

This report outlines the knowledge transfer process for the HPAC Platform, which is in place to facilitate the sharing of this knowledge, e.g. technology solutions developed based on HBP requirements, with external stakeholders, i.e. with developers of relevant HPC technologies and architectures. It largely centres on a collection of benchmarks, which reflects the needs of the HBP community for HPC- and Cloud-based services. A portfolio comprising nine benchmarks has been developed and implemented in a joint effort with the HBP ICEI project. The benchmarks are based on applications frequently used in the HBP, e.g. Arbor, NEST, NEURON and Elephant.

3.2 Validation and Impact

The outputs related to KR7.1 that have been produced so far mainly serve to help improve and augment the HPAC Platform services portfolio.

3.2.1 Actual Use of Output(s) / Exploitation

The **collection of science and use case documents** is used inside the Project for specifying the HPAC Platform architecture and the requirements it has to comply with.

The process defined in the **"Approach for platform validation for SGA2 defined" Report** is used to define functional tests and benchmarks for validating the HPAC Platform.

The process defined in the **"Knowledge transfer process for SGA2 defined" Report** is used to produce benchmark results. They are currently used internally, in preparation for the ICEI procurements.

3.2.2 Potential Use of Output(s)

The benchmark results, which constitute an outcome of the process defined in **"Knowledge transfer process for SGA2 defined" Report**, will be documented and shared with HPC and Cloud solution providers.

3.2.3 Publications

No publications yet.

3.2.4 Measures to Increase Impact of Output(s): Dissemination

The outputs created during the first year of SGA2 serve primarily internal purposes. In the context of knowledge transfer, outputs are planned that will need dissemination among providers of HPC- and Cloud-based solutions. We plan to use bilateral contacts with such solution providers, as well as relevant annual events like the ISC High Performance Conference, the Supercomputing Conference, as well as the EuroHPC Summit Week, to address a broader range of solution providers. Results will furthermore be provided to relevant technology projects, most prominently the European Processor Initiative (EPI).

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

This section reports the output produced by WP7.2, which has the goal to deliver data federation and data-intensive computing services and technology that will be used in the HBP and also by other Fenix user communities, as some of these technologies are developed for the Fenix infrastructure.

4.1 Outputs

4.1.1 Overview of Outputs

Table 2: Overview of outputs related to Key Result KR7.2

Output	Component number(s)	Component name(s)	Additional information
HBP/Fenix data transfer service	C416	Transfer services	Related milestone: MS7.2.1 (achieved in M11)
Private Cloud On a Compute Cluster (PCOCC)	C2825	User-managed virtual cluster on an HPC system through SLURM	

4.1.2 HBP/Fenix data transfer service

The "HBP/Fenix data transfer service" is a data transfer service that exploits the UNICORE FTP technology, combined with multiple other software components, to allow users to transfer data between the supercomputing centres in the HPAC Platform and Fenix. Its development has been partially funded by the HBP within Task T7.2.2 "Data transfer services enabled", led by Michele CARPENÉ (CINECA).

This data transfer service allows HBP users to move datasets between the HPC sites, both in "client to server" and "server to server" modes (see Figure 2). The technology adopted within the HBP to implement the Fenix data transfer service is the UNICORE FTP (UFTP), a high-performance data transfer service based on FTP. UFTP is also being considered as an option for implementing the Data Mover service in the ICEI project for the Fenix infrastructure. This service stages data between the archival and the active data repositories inside an HPC centre, whereas the data transfer services takes care of transfers between centres.

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

1) 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 transfer (server to server).
2) Jupyter notebooks available in the HBP Collaboratory:
   - HBP users can use the UFTP server via the UNICORE REST API from the Collaboratory.
   - The Collaboratory enables users to integrate data transfer steps inside more complex workflows.

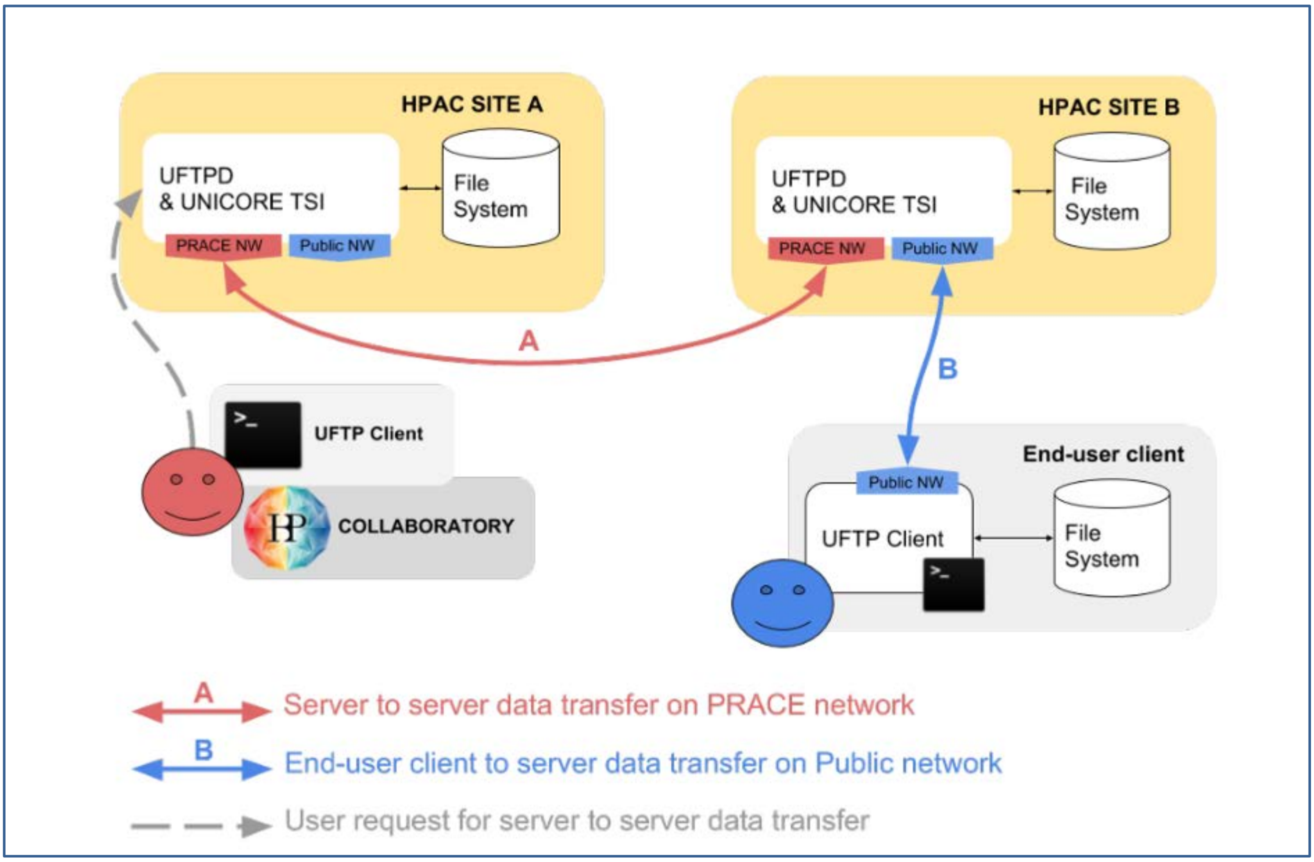

Figure 2: General view of the data transfer services in the HPAC Platform and Fenix

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.

The major elements enabling the data transfer services are the following:

pyunicore (Python library for using the UNICORE[5] REST API)

HBP users currently use this library to transfer data across multiple sites from the Collaboratory. The library has also been tested using a Jupyter notebook integrated in the Collaboratory.

Releases (Since May 2018):

- pyunicore 0.5.6: https://pypi.org/project/pyunicore/0.5.6/
- pyunicore 0.5.7: https://pypi.org/project/pyunicore/0.5.7/

UNICORE REST APIs

The REST APIs are essential for interacting with the UNICORE system, to transfer data and submit jobs within the HPAC/Fenix federation. The APIs have been tested using a Jupyter notebook in the Collaboratory.

Releases (Since May 2018):

- https://sourceforge.net/p/unicore/wiki/REST_API/

UFTP standalone client

This client can be directly used by HBP users from their own workstation to interact with the data transfer service. Through this client, it is possible to transfer large files from the user's workstation to any of the HPC centres (BSC, CEA, CINECA, ETHZ-CSCS, JUELICH-JSC), or directly between two of the centres, ensuring high performance and reliability.

[5] https://www.unicore.eu/

Additional information: https://unicore-dev.zam.kfa-juelich.de/documentation/uftpclient-0.8.0/

Releases (Since May 2018):

- UFTP client v0.80: https://sourceforge.net/projects/unicore/files/Clients/UFTP-Client/0.8.0/

Performance and functionality tests have been conducted for both “client to server” and “server to server” configurations over the public Internet and via the secure PRACE network, using the UNICORE UFTP standalone client. All tests have been performed using the last available UFTP version (v2.6.1).

Over the public network, the measured transfer rate was about **80/100 MB/s**. Using the PRACE network, the measured transfer rate was about **150 MB/s**. Data transfer tests over the PRACE network so far have involved CINECA, ETHZ-CSCS and JUELICH-JSC as target sites. A lot of work has been done to properly configure the UFTP endpoints on all the sites and some tuning has been required to complete the setup. While the initial testing was done by the above mentioned three centres, corresponding installations at BSC, CEA and KIT are already being worked on and will be in production soon. More details are provided in the internal report related to Milestone MS7.2.1, which can be made available on demand[6].

4.1.3 Private Cloud On a Compute Cluster (PCOCC)

PCOCC[7] (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 as part of Task T7.2.5 and will be made available in production for HBP users in the coming weeks after it has been fully tested and evaluated.

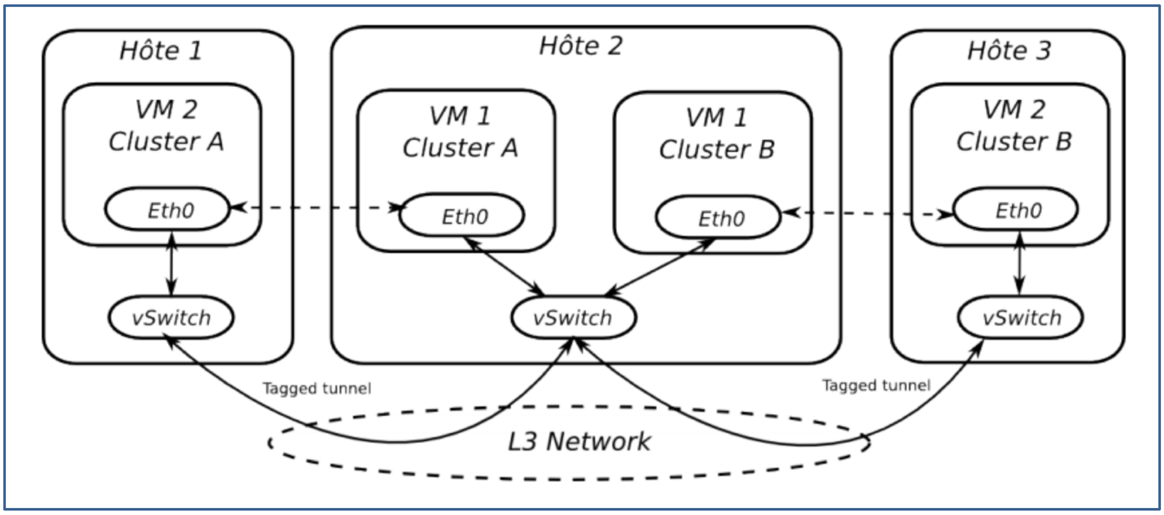

Figure 3: PCOCC general architecture

[6] Contact hpac-support@humanbrainproject.eu to receive a copy of this internal report.
[7] https://github.com/cea-hpc/pcocc

4.2 Validation and Impact

The UNICORE FTP (UFTP) service federation allows HBP and Fenix 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 Fenix e-infrastructure that the HPAC Platform builds on. The UFTP file transfer service can be used both with the command line client (UFTP standalone client) and from the Collaboratory (Jupyter notebooks), and it can be integrated with external services and tools. PCOCC will enable new utilisation models, as users will be able to run virtual clusters of several thousands of VMs, e.g. for running complex software stacks packaged in an image or for reproducing technical issues happening at large scale, without impacting production servers.

4.2.1 Actual Use of Output(s) / Exploitation

The data transfer service has already been successfully exploited 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 ability to perform data transfers over the PRACE network permits secure and reliable exchange of data among all connected sites, and avoids malicious access. No usage records for the PCOCC service exist yet.

4.2.2 Potential Use of Output(s)

Thanks to the data transfer service, users will be able to easily transfer data among all federated HPC sites, which enables the execution of new workflows, potentially distributed, involving different computing architectures and platforms, as well as long-term archiving of data. The introduction of the PCOCC will also reinforce this model and help users to execute deeply customised environments, using HPC resources through virtualisation.

4.2.3 Publications

The main publications for this KR were:

- Lena Oden, Christian Schiffer, Hanna Spitzler, Timo Dickscheid and Dirk Pleiter "IO Challenges for HumanBrain Atlasing using Deep Learning Methods - An in-Depth Analysis" 27th Euromicro Conference on Parallel, Distributed, and Network-Based Processing (PDP), Pavia 2019" *Accepted*
- Pol Santamaria, Lena Oden, Cesare Cugnasco, Yolanda Becerra, Raul Sirvent, Philipp Glock, Jordi Torres "Evaluating the benefits of Key-Value databases for scientific applications" *Under review*
- Lena Oden, Christian Schiffer, Hannah Spitzer, Timo Dickscheid und Dirk Pleiter, "Data staging techniques to resolve IO bottlenecks for Deep Learning on large scale brain images" *Under review*

In the context of Deep Learning for brain imaging, some I/O challenges have to be resolved. The first and third publication describe our analyses of the technical difficulties and solutions identified to overcome these. The second publication summarises an evaluation of key-value databases for scientific applications. The solutions found can by adopted by other research groups inside and outside the HBP for their projects; they build on technology available at the HPAC centres. All publications are the results of collaborations between domain and HPC experts.

4.2.4 Measures to Increase Impact of Output(s): Dissemination

Dissemination activities for the data transfer service included:

- HBP CodeJam #9, 26-28 November 2018, Palermo (Italy)

- The goal of the HBP CodeJam workshops is to catalyse open-source, collaborative software development in computational and systems neuroscience and neuroinformatics, by bringing together researchers, students and engineers to share ideas, present their work, and write code together.
- The UFTP tools were discussed with software developers to investigate possible improvements and relations with other projects or services, including functionality and performance improvements.

- 1st HPAC Platform Training, 11-12 December 2018, Barcelona (Spain)
  - Including presentation of the UFTP service to users and researchers, as well as main UNICORE software tools (e.g. the ucc client) and the integration with the HBP AAI. **Output:** HBP users learned how to use basic UFTP functionalities.

No dissemination activities took place yet related to PCOCC.

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

5.1 Outputs

5.1.1 Overview of Outputs

Table 3: Overview of outputs related to Key Result KR7.3

Output	Component number(s)	Component name(s)	Additional information
NEST (incl. NESTML)	C209, C2696, C510, 2679, C2678, C2789, C661, C662	C209: NEST - The Neural Simulation Tool C2696: Prototype implementation of framework and numerics for phenomenological neuron model with compartmentalized third factor plasticity C510: Continuous dynamics code in NEST C2679: Prototype NEST code with abstracted synapse model representations C2678: Prototype implementation of unified connectivity generation scheme for very large Networks C2789: NEST user-level documentation C661: NEST Support for Providers C662: NEST Requirements Management	T7.3.1: C209, C2696, C510 T7.3.4: C209, C2679 T7.3.5: C209, C2678 T7.4.5: C209, C2789, C661, C662
Arbor	C2689, C2690, C2691	C2689: Publication on Arbor methods and implementation C2690: HPC multi-compartment benchmark and validation suite C2691: Arbor library and frameworks (SGA2 release)	T7.3.2, T7.3.6, T7.3.7: C2689, C2690, C2691
TVB	C3030	C3030: HPC-for-TVB software package	T7.3.10: C3030
Multi-simulator multi-scale interaction	C2692, C2693, C2694, C2695	C2692: NeuGen: Interactive neuronal network geometry generation tool C2693: Simulations with NEST – NEURON coupling using MUSIC C2694: Development of a model with full spatial resolution for the extracellular potential activity using ug4	T7.3.3: C2692, C2693, C2694, C2695

		C2695: Simulations showing extracellular potential activity as result of an active network which couple the model with full spatial resolution (implemented in ug4) with MUSIC	

5.1.2 NEST

In the first year of SGA2, we released NEST 2.16.0 (C209). This brings a new solution for the storage of connections, which reduces memory consumption and network build times (concepts and performance data described in Jordan *et al.* 2018), and extends the scalability of the simulator into the exascale. The functionality has been considerably enhanced, including adapting the rate model framework for novel use cases (C510). The software framework for synaptic plasticity was extended to prepare for the representation of third factor plasticity and a corresponding pull request including a simple example (Clopath plasticity) made against NEST (C2696, C662; MS7.3.4). Unified representation of network structure, with and without spatial structure, and significantly improved threading performance for connecting complex networks was implemented as prototype; the user interface for unified connectivity generation scheme was drafted (C2678).

The corresponding high-level modelling language, NESTML, has been advanced to allow the expression of synapse as well as neuron models (C2679; MS7.3.2). We have also developed language concepts that permit synapse variables to be held in common over all instantiations (homogeneous case), thus allowing the memory consumption to be reduced to the minimum required for the scientific question at hand. On 21-22 November 2018, a workshop “Are we building the right user-level documentation” was held with the developers of neuromorphic hardware (SP9) at the EITN. The workshop results are the foundation for the concept of user-level documentation planned to be created in SGA3 (C2789; MS7.4.9). The work on documentation included reorganising the documentation in the NEST repository, and integrating documentation-specific technology into the NEST infrastructure, i.e. Sphinx as the documentation generator and Read the Docs as the platform for publishing documentation online. The purpose of this was to improve discoverability and usability of the documentation for researchers as outlined in the discussions with users and developers. Extensive notes were produced for each release and communication with the simulation software maintainers of HPAC and NRP continued (C661).

5.1.3 Arbor

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.

The key achievements in the past year have been the delivery of an open source benchmarking and validation suite (see Deliverable D7.3.1 for details) and the development of the core simulation library and framework. The benchmarking and validation suite (C2690) is available online www.github.com/arbor-sim/nsuite and at the time of release provides:

- Automatic download and compilation of Arbor, NEURON and CoreNeuron simulation engines on HPAC systems
- A framework for defining and running performance benchmarks and validation tests
- Two benchmarks and one validation test
- Support for using continuous integration on Piz Daint at ETHZ-CSCS

Released in M12, the framework will continuously test and benchmark HBP simulation engines, with new benchmarks and validation tests to be added throughout the second year of SGA2. A key finding from the benchmark suite is that Arbor requires significantly less HPC resources, both time and memory, than equivalent simulation software (NEURON and CoreNeuron). When integrated into HBP workflows, Arbor will deliver performance and model scales on HBP infrastructure beyond the current state of the art for multi-compartment models.

The development of the core Arbor library (C2691) saw the release of version v0.1 and v0.2 during the last year. Significant features introduced in this period include: electrical gap junctions, back end optimisations (Intel SIMD code generation, a highly-parallel GPU matrix solver), flexible runtime management of HPC resources (GPUs, thread, MPI) and code maintenance.

5.1.4 TVB

We have concentrated on the design of the Domain Specific Language (DSL) for neural mass models which TVB-HPC will use in order to generate HPC kernels. The initial design has been published and the code base is available in the GitHub repository (https://github.com/the-virtual-brain/tvb-hpc/; C3030). This work has been done in close collaboration with Marmaduke Woodman and Jan Fousek (AMU).

5.1.5 Multi-simulator multi-scale interaction

NeuGen, the network generation tool (C2692) has been coupled to VRL Studio to provide a convenient user-interface. With the help of NeuGen, large realistic networks can be generated based on stochastic parameters and can be exported to suitable output formats (hoc, NeuroML, ugx). Furthermore, FG has started code optimisation of NeuGen for parallelisation.

As a proof of concept, a simple simulation with NEST and NEURON has been established, where simple neurons in NEST provide synaptic input for NEURON (C2693). The NEURON simulation comprises a pyramidal cell with Hodgkin Huxley currents receiving input in form of time stamps at alpha synapses. Data transfer was established using MUSIC (C347). This model has been extended with the fully resolved model in UG4 (C2695). A physically realistic fully resolved model (3D resolution; C2694) based on Maxwell's laws and ohmic currents has been derived and implemented. It will now be tested and optimised.

5.2 Validation and Impact

The improvements of the NEST simulator now allow it to be used on the largest HPC facilities, enabling simulations of larger network models in shorter time (Jordan *et al.* 2018), while the increased functionality in the area of synaptic plasticity widens the applicability of NEST in studies of network function. NESTML allows the user to define synapse models without any knowledge of C++ or particular software development expertise. This enables neuroscientists without computational background to use the simulation framework, and thus contributes to the sustainability of NEST by considerably reducing the number of code lines. Substantial impact was achieved by a joint study (van Albada *et al.* 2018) with neuromorphic computing experts demonstrating the simulation of a full-scale cortical microcircuit model on the SpiNNaker hardware. This is a breakthrough, as all larger networks are less densely connected. The results were validated by high-precision simulations with NEST.

Feedback from the community inside and outside the HBP has been gathered in workshops like "Are we building the right user-level documentation" to validate the present release and shape the future research programme.

The present release of Arbor enables much faster simulations of multi-compartment model neurons than with all previously available codes. This is documented in a preprint (https://arxiv.org/pdf/1901.07454.pdf) shortly to be published as conference proceedings.

The conceptual work with regards to TVB carried out in the first half of SGA2 resulted in design of the implementation work to be carried out in the second half.

The work related to multi-simulator, multi-scale interaction enables the reliable coupling of multiple simulators, bridging the scales between time-based problems, 1D cable models and fully resolved 3D models. The fully resolved model, which can additionally represent ephaptic signals, is currently being tested and optimised. With the network generation tool, we provide the networks for scientific computation, required by simulators like NEURON and UG4. NeuGen is a useful tool for easy and automatic access to large networks.

5.2.1 Actual Use of Output(s) / Exploitation

NEST is the main neural simulator in the HBP at the spiking neuron network level. In addition, it is used for the validation of the neuromorphic hardware systems. NESTML permits rapid development of neuron and synapse models that extend NEST's built-in selection, thus providing the simulator with a high flexibility. The NEST simulator is used for studies of network dynamics, for example, the recently released multi-area model of vision related areas (Schmidt *et al*. 2018) containing millions of neurons. However, NEST is also used for studies of network function and in the Neurorobotics Platform. We are aware of 33 peer-reviewed, published articles in M1-M12 using NEST with authors from inside and outside the HBP. At workshops and the annual NEST Conference, users and developers get into direct contact and the resulting insights shape the future NEST development. Furthermore, feedback is gathered through the openly accessible issue-tracking on GitHub, the NEST user mailing list and in fortnightly open video conferences.

Arbor provided benchmarks to the ICEI benchmark suite that will be used for testing and procurement of computing systems for the Fenix infrastructure. Arbor is the only simulation engine in the benchmark suite that provided benchmarks for GPU-accelerated systems. The benchmark suite had also been suggested by the HBP reviewers.

The TVB code is being tested by external users in applications which require extensive parameter exploration for subject fitting.

5.2.2 Potential Use of Output(s)

In the next and final phase of HBP (SGA3), the maturity of the NEST code needs to further increase; the outputs of SGA2 provide the necessary preparations for this. The developments in NESTML will allow a reduction of the code base of NEST in the next years, thus improving its maintainability. In addition, we have developed a concept for user-level documentation of NEST to be implemented in the future. This effort is coordinated with the development of user-level documentation for the neuromorphic hardware system, to ensure a homogeneous user experience and easy transfer of knowledge.

Arbor is new software, developed in SGA1 and SGA2. To date, the focus has been on developing the software framework and features needed by the user community. It should be emphasized that the software has a higher quality and performance than the existing simulation engine NEURON, with better internal and external APIs, cleaner code, and extensive unit testing. Results from the benchmarking suite show that Arbor has two main advantages over NEURON: it requires an order of magnitude less time and memory to run large models, and it runs efficiently on diverse HPC hardware (e.g. GPUs and KNL). Hence Arbor has great potential to run larger and more detailed models on the HBP HPC infrastructure than currently possible, and Task T7.3.7 is reaching out to users to establish collaborations.

Given the increased performance gain offered by the produced code for TVB, applications targeting large groups of subjects or extensive parameter exploration for complex neural models will become feasible. This will increase the amount of information for analysis and increase the reliability of observed simulated behaviour and of the conclusions drawn from these observations.

To easily obtain realistic neuronal data, NeuGen is a fast solution to provide realistic morphologies of single neural cells and huge networks of neural cells. Coupling NeuGen to VRL Studio yields a great improvement in user-friendliness and also provides the possibility to directly communicate with simulators, which have been integrated in VRL-Studio as plugins. Furthermore, the fully resolved model is to our knowledge a unique model representation in three space dimensions, depicting ephaptic coupling, and it can extend lower dimensional neuronal simulators such as NEURON or NEST by this functionality.

5.2.3 Publications

The main publications for this KR were:

There was one peer-reviewed publication (P1456) on NESTML in the first year of SGA2: Blundell *et al.* 2018, Code Generation in Computational Neuroscience: A Review of Tools and Techniques, *Frontiers in Neuroinformatics* 12 (2018): 68. The publication explains the general strategy for automatic code generation in TVB-HPC and generic model descriptions. (https://www.frontiersin.org/articles/10.3389/fninf.2018.00068/full).

In February 2018, a peer-reviewed article (P884) on NEST development (Jordan *et al.* 2018, Extremely Scalable Spiking Neuronal Network Simulation Code: From Laptops to Exascale Computers, Frontiers in Neuroinformatics 12:2; https://www.frontiersin.org/articles/10.3389/fninf.2018.00002/full) appeared, which has not been reported in SGA1. This article documents the latest simulation kernel (5g) available in the present NEST release. The article demonstrates that memory consumption of an individual compute node is independent of total size of the neuronal network.

A paper about Arbor's design, implementation and benchmark performance, was presented at the PDP 2019 conference in Pavia, Italy, in M11 (C2689). The paper was titled: "Arbor - A Morphologically-Detailed Neural Network Simulation Library for Contemporary High-Performance Computing Architectures". A preprint is available online (https://arxiv.org/pdf/1901.07454.pdf), and an author's copy will be provided in the conference proceedings.

5.2.4 Measures to Increase Impact of Output(s): Dissemination

Community-wide interest in the NEST Simulation framework and related projects is annually expressed at the NEST Conference. The about 40-50 participants, users and developers from over ten countries (2018), come together to discuss latest developments, related projects and the future of the simulator. They took part in lively interactions during co-organised satellite activities such as workshops and hackathons, thus enhancing knowledge transfer as well as international collaborations. The aforementioned workshop at the EITN was held with the developers of neuromorphic hardware (SP9) aiming at users of NEST, SpiNNaker and BrainScaleS.

The publication van Albada *et al.* 2018 (Frontiers in Neuroscience; https://www.frontiersin.org/articles/10.3389/fnins.2018.00291/full) on the SpiNNaker-NEST comparison of the cortical microcircuit model of Potjans & Diesmann (Cereb Cortex, 2014), which has been supported by SP7, has received extensive media attention (e.g. Frontiers Communications, TOP500, DNA India, Science Friday radio, Science Node magazine, Forschungszentrum Jülich and University of Manchester websites, on best paper selection page) increasing the visibility of NEST.

An important dissemination event was the 1st HPAC Platform Training (https://www.bsc.es/education/training/other-training/1st-hpac-platform-training-event) held at BSC on 11-12 December 2018. NEST and Arbor tutorials were presented on how to run them on HPC systems provided by the HPAC Platform. The Arbor tutorial taught using the Python front end on Jupyter notebooks directly on CSCS GPU-accelerated supercomputer Piz Daint.

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

SGA1 Deliverable D7.6.3 provided an overview of the different visualisation activities in the HPAC Platform; a summary is also available in the HPAC Guidebook[8]. The focus in the last year was on software for in situ visual analysis of simulation results, including a pipeline that connects simulator and analysis tools, and on implementing a back end that enables visualisation and analysis tools to run on the new Fenix infrastructure. A new visualisation software catalogue is under development that will provide an overview of which tools are available for which purposes.

6.1 Outputs

6.1.1 *Overview of Outputs*

Table 4: Overview of outputs related to Key Result KR7.4

Output	Component number(s)	Component name(s)	Additional information
In situ Pipeline	C2680, C2697	C2680: HPAC integration of in situ pipeline, C2697: Flexible in situ pipeline coupling NEST to visual analysis tools	
Deployment and extension of visualisation tools at CINECA	C2681	Development, extension and deployment of visualisation tools for HPAC	
Release of a prototype for remote visualisation and interaction	C2682	Remote high-fidelity visualisation and interaction middleware	
Hands-on workshop and roadmap definition	C2698	Evolution of visual analysis tools towards the in-situ pipeline	
Visualisation software catalogue – design and first version	C2707	Semantic Visualisation Software Catalogue	

6.1.2 *In Situ Pipeline*

In situ visualisation addresses the problem of flexible, visual analysis of very large data generated during the execution of simulations, emerging from simulators developed and supported in the HBP at scale. It enables a direct view into a running simulation and thereby makes error detection possible and more convenient if steering of the simulator is additionally enabled. Using flexible software interfaces, various types of existing and new visual analysis tools can be connected through different network protocols (C2698).

The *in situ* pipeline is implemented in two separate libraries called nesci (a conan interface for neuronal simulators) and contra (transport conduit data between processes), which are both

[8] https://hbp-hpc-platform.fz-juelich.de/?page_id=1800

available in public git repositories (https://devhub.vr.rwth-aachen.de/VR-Group/nesci and https://devhub.vr.rwth-aachen.de/VR-Group/contra, respectively). These repositories also include documentation on how the libraries are intended to be used. In the current state, the pipeline has been successfully connected to NEST and Arbor, whereas the support for TVB is currently under development. The integration into NEST is a key aspect of our "Flexible *in situ* pipeline coupling NEST to visual analysis tools" (C2697) and is realized by providing a nestio-based streaming module in the public repository (https://devhub.vr.rwth-aachen.de/VR-Group/nest-streaming-module). For Arbor, a proof-of-concept integration can be found in the risotto branch in the official repository located at https://github.com/arbor-sim/arbor.

The other aspect of our work on the *in situ* pipeline (see C2697 and especially C2680) focuses on the support for and the deployment on the HPAC Platform. In order to achieve this, we implemented a flexible layer of network protocols to address the specific needs of such HPC environments. We plan to deploy the *in situ* pipeline to the HPC systems in the second year of SGA2.

6.1.3 Deployment and extension of visualisation tools at CINECA

In collaboration with Prof. Michele Migliore's group, the visualisation tools ViSimpl and NeuroTessMesh were compiled and deployed on GALILEO, a supercomputer at CINECA, and different features added (Task T7.3.8; see Figure 4).

ViSimpl (C2681) provides 3D particle-based rendering that allows simulation data to be visualised with their associated spatial and temporal information for enhancing the knowledge extraction process. It provides abstract representations of the time-varying magnitudes supporting different data aggregation and disaggregation operations and giving focus and context clues. In addition, ViSimpl tools provide synchronised playback control for the simulation being analysed and allows to perform selection and filtering operations that rely on NeuroScheme[9]. All these views are loosely coupled and can work both, in centralised and distributed computing environments.

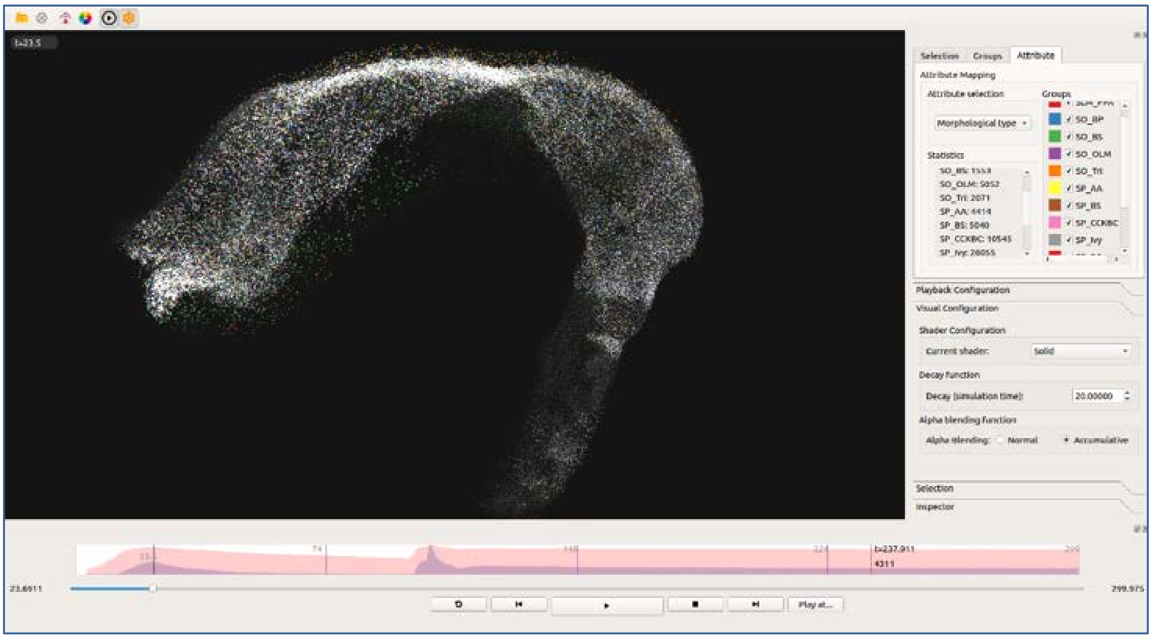

Figure 4: Simulation of the Hippocampus in ViSimpl with selection of different groups of neurons

Courtesy of Michele Migliore, on behalf of the Hippocampus Team of SP6 (EPFL, CNR, IEM HAS, UCL)

[9] https://github.com/gmrvvis/NeuroScheme

In terms of the mentioned extensions, ViSimpl now allows to discriminate the activity of different groups of neurons visually as well as to present detailed information of a neuron when clicking on it. Some other extensions have been added to enhance the interaction with the data.

NeuroTessMesh (C2681; see Figure 5) is a tool specifically designed to visualise neural circuits comprising large numbers of cells. It facilitates the recovery and visualisation of the 3D geometry of cells included in databases, such as NeuroMorpho, and provides the tools needed to approximate missing information such as the soma's morphology. NeuroTessMesh takes as input the available compact, yet incomplete, morphological tracings of the cells. It uses a multi-resolution approach that combines an initial, coarse mesh generation with subsequent on-the-fly adaptive mesh refinement stages using tessellation shaders. The adaptive refinement process performed on the graphics card generates meshes with geometries of good visual quality at a reasonable computational cost, both in terms of memory and rendering time.

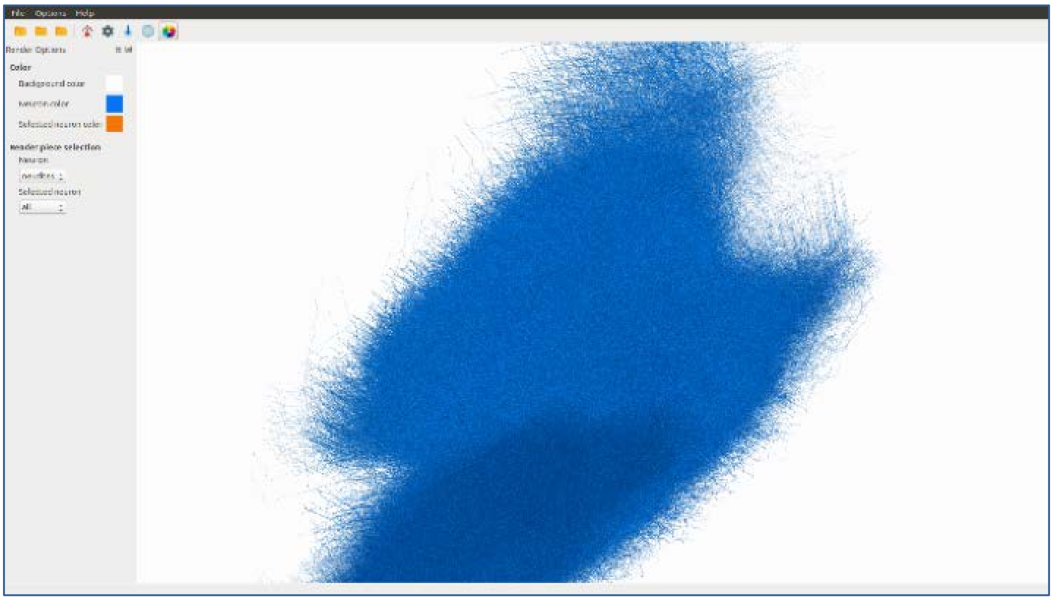

Figure 5: Morphologies of a region from the hippocampus viewed in NeuroTessMesh

Courtesy of Michele Migliore, on behalf of the Hippocampus Team of SP6 (EPFL, CNR, IEM HAS, UCL)

With the newly developed extensions, NeuroTessMesh is now able to assign alternative visual codification to different parts of a morphology. Additionally, NeuroTessMesh can deal with occlusions and it can also provide information about neuron activity. Apart from that, in order to load huge datasets an OpenMP-based optimisation has been included. Such large datasets are now accessible after the deployment of the tool on GALILEO, one of the HPC systems at CINECA.

6.1.4 Release of a prototype for remote visualisation and interaction

The remote visualisation back end (RVB) service under development (C2682) is planned to provide interactive remote rendering capabilities, which allow users to establish visualisation sessions on the HBP infrastructure and to manage computing sessions for running simulators within *in-situ* visualisation scenarios. The RVB will directly stream the visualisation results to a standard web browser, providing visual analysis capabilities for very large data anytime and anywhere, using the HPAC Platform and potentially Fenix resources as compute back end. As a main goal, the RVB service will provide expert and non-expert users with easy accesses to HPAC resources for launching applications and visualising results in an easy and uniform manner. The integration of the client in the Collaboratory will make these capabilities available to all HBP users.

The developed framework is composed of a library (ReMo) and a set of services (ReMoToo, ReMoLON-Frontend, ReMoLON-InstanceManager), built on top of the former. The framework follows the scheme depicted in Figure 6 to allow the users to establish interactive remote visualisation sessions. ReMo (https://github.com/gmrvvis/ReMo) is a low-level library for managing video streams and connecting video generators with video consumers. It uses the "webstreamer" library (https://github.com/HBPVIS/webstreamer) as the transport layer developed by RWTH (in kind). ReMoToo (https://github.com/gmrvvis/ReMoToo) is the service running on the visualisation nodes. It relies on ReMo+webstreamer capabilities for capturing the whole desktop, applying compression and streaming it to the client, while receiving user interaction events. Besides, ReMo provides features for enhancing the interaction, such as clipboard management or layout-independent keyboard interaction. ReMoLON (https://github.com/gmrvvis/ReMoLON) is composed of two services: the front-end and the "instance manager". The former provides a web server and acts as the entry point for authentication, session managing and resource requests. The latter runs as a service on the visualisation nodes and it is in charge of allocating resources needed and creating ReMoToo instances, as required.

The developed remote rendering back end is a custom-made middleware that allows the running of any application that has been built on the machines where the back end is deployed. The entry point for the back end is a web page that allows users to create and manage sessions, enabling them to run potentially any application from their web browser. The integration of the back end in the Collaboratory is currently ongoing.

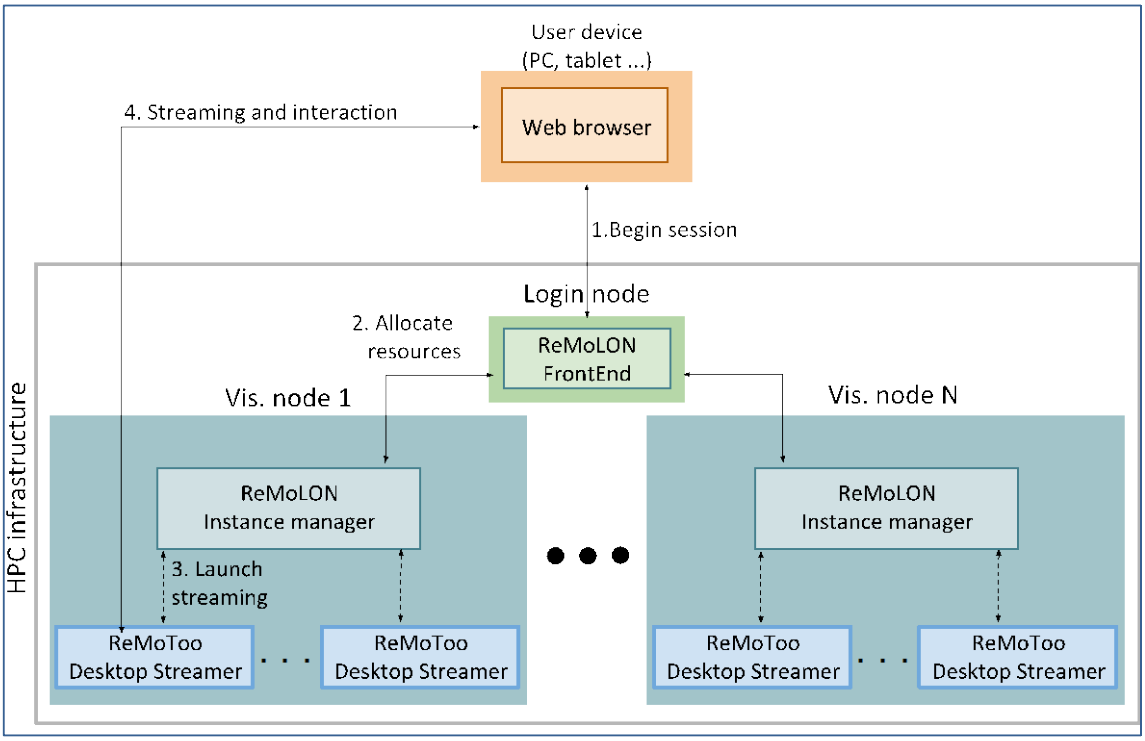

Figure 6: Scheme of components of the remote rendering framework

6.1.5 *Hands-on workshop and roadmap definition*

On 19-21 September 2018, the RWTH, UPM and URJC groups participated in a hands-on workshop to study different proofs of concept of the planned developments and define a roadmap for Task T7.3.8. In addition to the discussions about the roadmap, a significant effort has been done in the evolution of NeuroScheme (ConGen; C2698), a tool for creating and editing networks using different levels of abstraction built on top of the NeuroScheme's framework developed during RUP and SGA1. The output of this tool can be used to feed simulators and plays a key role in closed-loop in-situ visualisation workflows. NeuroScheme (ConGen) allows to define population of cells, to visually establish connections among them, to define the properties of both cells and connections, as well as to create and connect them to the population.

The roadmap specifies the next steps required to integrate the *in-situ* pipeline developed at RWTH into the visualisation tools developed at URJC. First, the focus was set on integrating the pipeline into a single tool (ViSimpl, C2681), thereby enabling its coupling with running NEST simulations. More specifically, the focus will be on the identification of necessary features and API extensions on both sides. Therefore, the roadmap includes a research stay at URJC by Simon OEHRL, developer of the pipeline at RWTH, to execute the planned implementation work in close collaboration with the ViSimpl developers. The research stay is also meant to allow him to detect incompatibilities and to adjust the tools accordingly. Features necessary for the start of the integration will be executed beforehand. A final working implementation should be in place by the end of 2019 (M21). This work ultimately aims to implement a use case including NeuroScheme (ConGen) and ViSimpl, which is coupled to NEST, for fast, iterative development of large neural network models simulated in NEST. Due to its interface definition, the pipeline also provides easy means to couple running simulations to non-visual analysis tools such as Elephant. In SGA2, however, the focus is only on coupling running simulations to visual analysis tools that are specifically identified in the work plan.

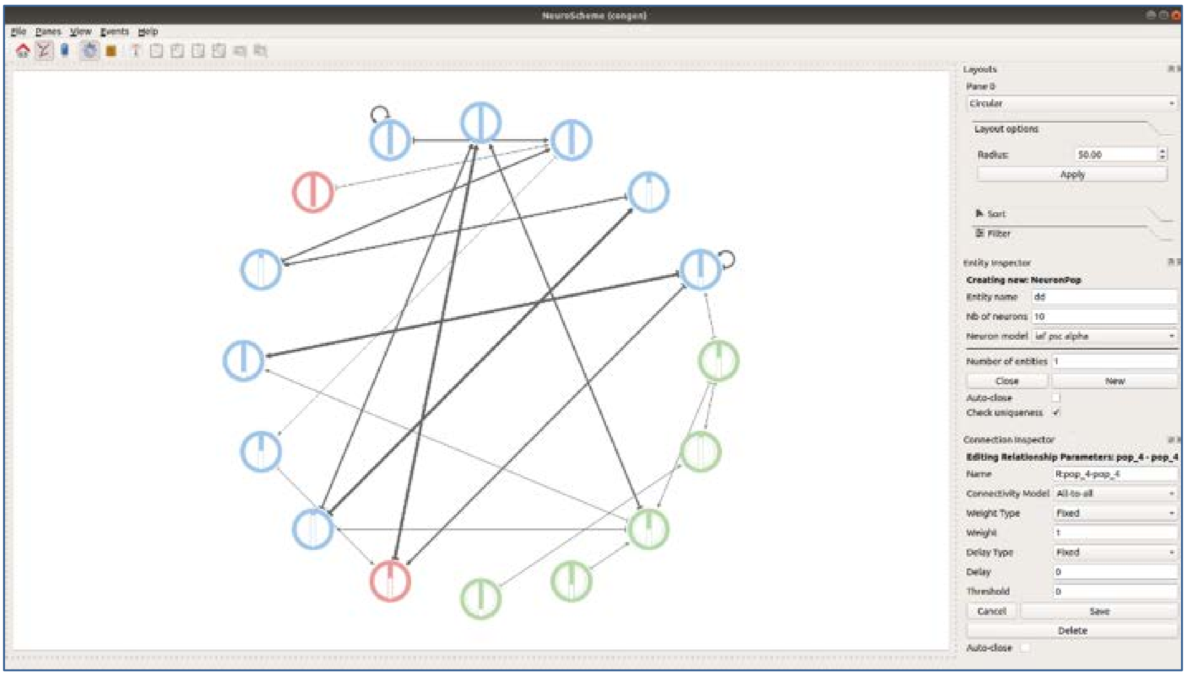

Figure 7: Circular layout of a network created using ConGen

6.1.6 Visualisation Software Catalogue

The visualisation software catalogue (C2707) will provide the HBP's user community and external users with information and support on visual analysis tools and software for neuroscientific data analysis. Therefore, the catalogue will integrate metadata on visualisation applications developed within the HBP (such as ViSimpl) and information about externally developed tools (such as VTK), so that scientists can get a broad view of tools and software supporting their specific needs in terms of visual analysis.

In the first year of SGA2, the work focused on researching established standards for describing visualisation, visualisation software, as well as software in general. Various types of meta data schemas have been identified and ontologies considered for future use on the catalogue. An overall system design has been developed and prototypically implemented. This includes the development of a meta data schema for visualisation software. For the first online version of the visualisation software catalogue, a BaseX database has been established, which will be publicly accessible in March 2019 (M12) through a web interface as planned.

During the conceptualisation phase of the visualisation catalogue, it became evident that the infrastructure, as designed and implemented in prototype form, showed a significant overlap with the functionality and structure of the HBP Knowledge Graph (KG). Therefore, in the second year of SGA2, it is planned firstly to integrate the developed metadata schema into the KG, and then define a metadata curation process for integrating metadata into the KG. To broaden the scope of this activity, it has been decided to consider not only visualisation software, but also general software tools developed within the HBP and externally. This will enable the integration of metadata for software into metadata for datasets via the established curation processes (e.g. to annotate datasets with metadata for software used to process this data). In addition, this makes it possible to use software metadata for supporting provenance data in the KG besides the original plan to support users by means of semantical search of (visualisation) software through a web interface (which will be implemented by the KG web front end).

To streamline this process, UT, as main developer of the visualisation catalogue, established a close collaboration with the HBP Knowledge Graph team, the data curation team, and the Collaboratory developers, who also maintain the HBP-internal software catalogue. It is planned to integrate this information into the KG via the software metadata curation process described above. In this way, a potential duplication of efforts can be avoided, and we can ensure that the activities are well integrated and aligned. However, UT still plans to maintain the initial release of the catalogue in the HPAC Guidebook via automated scripts, which extracts software metadata from the KG and generates an HTML-based version of the Guidebook. Thus, no extra work for updating both sources will be necessary and inconsistencies will be prevented.

6.2 Validation and Impact

6.2.1 Actual Use of Output(s) / Exploitation

The *in situ* pipeline (C2679, C2680) is used to enable visual analysis of neuronal network simulations with NEST and Arbor during runtime.

The visualisation tools (C2681) are being used by the groups of Michele Migliore (CNR), Idan Segev (HUJI), Egidio Di Angelo (UNIPV), Eduardo Ros (UGR), the SimLab Neuroscience (JUELICH-JSC), and Simo Vanni (University of Helsinki and Helsinki University Hospital). Ongoing work on the visual analysis of the hippocampus model from the Migliore group (Brain Simulation Platform; SP6 and CDP4) is in an early stage, since the tools will have to be adapted to the needs of this researcher group.

With regards to remote high-fidelity visualisation and interaction middleware (C2682), a first prototype of the remote rendering back end is available (MS7.3.13).

ConGen (C2698) is used by the SimLab Neuroscience (JUELICH-JSC) for the creation, editing and connection of networks, although it is currently in an early development stage.

6.2.2 Potential Use of Output(s)

A potential use of the *in situ* pipeline (C2679, C2680) is to realise multi-simulator scenarios to simulate heterogeneous neural models connecting NEST, Arbor and TVB. In addition, it can also be used to implement brain plasticity by adding steering to the pipeline. The *in situ* pipeline has reached TRL 4-5.

ReMo, ReMoToo and ReMoLON (C2682) are at TRL 4 and could be used in many use cases where visualisation is needed, inside the HBP and even outside the HBP, not only in the neuroscience domain.

ViSimpl and NeuroTessMesh (C2681) have TRL 4 and could be exploited by other HBP groups that need to do visual analysis and data exploration, in addition to those mentioned above. Many of the techniques developed are domain-independent, so they are suitable to be used also for other visualisation problems outside the HBP.

ConGen (C2698) is at TRL 4 and could also be exploited by other HBP groups besides the one mentioned above. The techniques developed are domain-independent, so they are suitable to be used on different types of networks and populations also outside the HBP.

The visualisation software catalogue (C2707) is at an early stage with a first proof of concept available (TRL 3). Focus lies on the integration into the Knowledge Graph and the fusion with the currently existing software catalogue in the Collaboratory.

6.2.3 Publications

- Simon Oehrl, Jan Müller, Jan Schnathmeier, Jochen Martin Eppler, Alexander Peyser, Hans Ekkehard Plesser, Benjamin Weyers, Bernd Hentschel, Torsten W. Kuhlen, Tom Vierjahn (2018) Streaming Live Neuronal Simulation Data into Visualization and Analysis. In: Yokota R., Weiland M., Shalf J., Alam S. (eds) High Performance Computing. ISC High Performance 2018. Lecture Notes in Computer Science, vol 11203, p. 258-272, Springer, Cham. doi: 10.1007/978-3-030-02465-9_18
    - The above publication (P1759) outlines how data can be streamed from neuronal simulations to visualisation and analysis tools, which is an important prerequisite for several use cases in the context of interactive (super-) computing, e.g. for enabling steering of simulations. Being able to visualise or analyse a simulation live allows modification of simulation parameters or stopping the simulation if the results are not as expected, which in turn saves limited compute resources.

6.2.4 Measures to Increase Impact of Output(s): Dissemination

"New approaches on Visualization in Neuroscience", Óscar David Robles Sánchez, 29 June 2018. This presentation was made for the CEIG 2018 (Computer Graphics Spanish Conference), the annual meeting of researchers from Computer Graphics and Visualisation (covers C2681 and C2698).

To emphasize the development of the visualisation catalogue (C2707), as well as a community around visual analysis tools for neuroscience, a fortnightly open video conference on visualisation was established in November 2018.

The visualisation catalogue team participated in the 1st HPAC Platform Training in December 2018 in Barcelona to present the catalogue and visualisation methods in general for use in neuroscience. The next participation will be in a workshop in Jülich in June 2019. Additionally, the team supported Sarah Haas (JUELICH-INM1) and Timo Dickscheid (JUELICH-INM1) in identifying visualisation software for interactive annotation in 3D upon request.

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. Some development work was done during the first twelve months of SGA2.

The biggest challenge faced by the HPAC Platform has been providing seamless access to HPC data from Jupyter notebooks within the Collaboratory. Behind-the-scenes work is on-going on several options. Specifically, JUELICH-JSC and ETHZ-CSCS are working on OpenStack solutions using features of GPFS. Kubernetes deployments with closer access to the file systems or running a JupyterHub service closer to the HPC systems are also being considered. Work is on-going to crack this difficult problem.

The original idea of the weekly HBP Joint Infrastructure Coordination meetings was to use them as a way to introduce a more use-case driven approach; this was initially hindered by several factors. However, these meetings are now on a much better path and we are settling on a better format now 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.

One final point is that we are waiting on the finalised Architecture Specification document (D7.1.1) to help ensure we are augmenting the service offerings in the right way and in the correct direction.

7.1 Outputs

7.1.1 Overview of Outputs

Table 5: Overview of outputs related to Key Result KR7.5

Output	Component number(s)	Component name(s)	Additional information
Integration and Operation of Low-level Infrastructure	C263, C329, C409, C410, C416, C417, C571, C1092, C1094, C1096, C1097, C1112, C1129, C2828, C2831, C2832, C2834, C2837-C2843, C2849, C2855, C2860, C2861, C2870, C2875	See Annex A: Additional information about releases and major updates	
Policy Management	C2871, C2872, C2873	C2871: HPAC Platform resource allocation policy summary C2872: HPAC Platform security policy review C2873: HPAC Platform Review of Data Management Plan and Data Policy Manual documents	
Infrastructure Services	C263, C329	C263: UNICORE C329: SP7 Federated HPAC Data and Computing Services	
Platform developer services	C329	C329: SP7 Federated HPAC Data and Computing Services	
Platform integration services	C263	C263: UNICORE	
Performance optimisation	C2709, C2710	C2709: Intermediate code optimisation report C2710: Final code optimisation report	
User support and documentation	C337, C1114	C337: HPAC Platform Guidebook for SGA2 C1114: HPAC Ticket System	
Education and training	C2867	C2867: HPAC Platform training programme	

7.1.2 Integration and Operation of Low-level Infrastructure

The operation of the low-level infrastructure at the partner sites, comprising the compute-, data- and network infrastructure services, has been coordinated. Day-to-day operational activities have been performed to ensure a continuously high availability of the low-level services. At ETHZ-CSCS, this encompassed also the support for the early approved ICEI projects and associated users.

A strategy for a GDPR-compliant account deletion process has been agreed and implemented, and the definition of the uptime key performance indicators (KPIs) for different low-level services has been synchronised. Work on the integration of the CEA infrastructure has been progressing. An important aspect was the analysis done for the integration of Kerberos authentication and authorisation with federated services. When finalised, the integration of the CEA infrastructure will be described in Deliverable D7.5.1, which is delayed due to unexpected technical issues to be solved. The issues identified and a tentative timeline are summarised in an internal, preliminary version of D7.5.1. The current estimation is that the CEA integration will be finalised in approximately M20.

At RWTH Aachen, a storage system as part of the on-site visualisation system has been integrated into the low-level infrastructure. Together with the platform integration by Task T7.5.5, this enables data transfer to Aachen within the HPAC Platform via UNICORE for the purpose of on-site interactive visualisation.

7.1.3 Policy Management

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. The output of the Data Governance Working Group (DGWG), namely the Data Management Plan (DMP), is also currently under review. Finally, the federated resource allocation policies used within ICEI are being assessed, since ICEI provides the agreed mechanisms for resource allocation within HBP.

7.1.4 Infrastructure Services

Regular maintenance of the various UNICORE installations has been performed including bug fixes for abnormal installation problems. Some additional issues have been addressed at some sites as part of the daily operations.

On the Neurorobotics Platform (NRP) side, ETHZ-CSCS has now fully integrated Piz Daint into the NRP workflow, meaning that NRP have completed building all the infrastructure needed in order to automatically launch NRP jobs on Piz Daint, on demand. This work has not yet been moved into full production. The next steps are to plan the work to allow the NRP to scale their brain simulations and this will set the work plan for the next six months.

A breakdown of the achievements are as follows:

1) Dynamically spawn single-container NRP on Piz Daint:
   a) Single-container back end
   b) GPU is used by Gazebo
   c) Front-end server runs on a Pollux VM (possibly with hardcoded parameters)
   d) Front-end and back-end servers communicate bi-directionally through an SSH tunnel setup by the back end upon spawning
   e) Selected single user execution (user bp000025) as proof of concept
   f) Automated launch of proxy using UNICORE over REST API
   g) Authentication done over OIDC Collaboratory server and authorisation token passed on to UNICORE by NRP
2) Currently in a testing phase is:
   a) Run parallel NEST within a single container
3) Next steps will include the ability to dynamically spawn NRP with parallel NEST on multiple nodes on Piz Daint:
   a) One back end is spawned
   b) Parallel NEST to be containerized
   c) CLE spawns parallel NEST via UNICORE
   d) Simulation parameters are passed to the NEST simulation at launch

e) Dependency on custom MPI code (modified MVAPICH2) is replaced by a standard MPICH3 implementation

f) Parallel NEST make use of Cray MPI (via Shifter MPI support)

7.1.5 Platform developer services

1) Deployment of visualisation tools for SP6

Brain Simulation Platform (SP6) researchers requested the installation of RTNeuron on CSCS' Piz Daint system to aid their workflows. RTNeuron[10] is software developed in the Blue Brain Project (BBP) for rendering and visualising neurons and the results of neuronal simulations. Since the installation turned out to be challenging, Brayns[11] (also developed in the BBP) has been installed instead, which is functionally equivalent to RTNeuron and is already used in production by them.

CSCS support staff contacted the Brayns developers, who suggested to use the Docker container image of the software instead of trying the native installation, due to the complexity of this task.

The Docker images for the Brayns visualiser[12] and the corresponding user interface[13] were retrieved, however the first attempts to run test cases bundled in the container images were unsuccessful. The outlook is to try again running the SP6 test case within the container on Piz Daint as soon as possible and keep communicating with the requestors through the HBP ticketing system.

2) Use of CSCS Continuous Integration (CI) services by the *Arbor* development team

The development team of the Arbor high-performance library for computational neuroscience simulations (see KR7.3 in Section 5.1.3) was given access to the CSCS Jenkins CI service during the HBP Summit 2018. The team has used the service to work on the delivery of a multi-compartment benchmark and validation suite on HPC systems (Deliverable D7.3.1). The suite supports compilation of the multi-compartment simulation engines Arbor, NEURON and CoreNeuron. Additionally, it features benchmarks and validation tests, which it can run in automated pipelines (such as CI servers). A proof of concept has been set up on the CSCS Jenkins CI service (see Figure 8). The team plans to continue to improve the suite and proof of concept over the rest of SGA2.

[10] https://developer.humanbrainproject.eu/docs/projects/RTNeuron/2.11/index.html
[11] https://github.com/BlueBrain/Brayns
[12] https://hub.docker.com/r/bluebrain/brayns
[13] https://hub.docker.com/r/bluebrain/brayns-ui

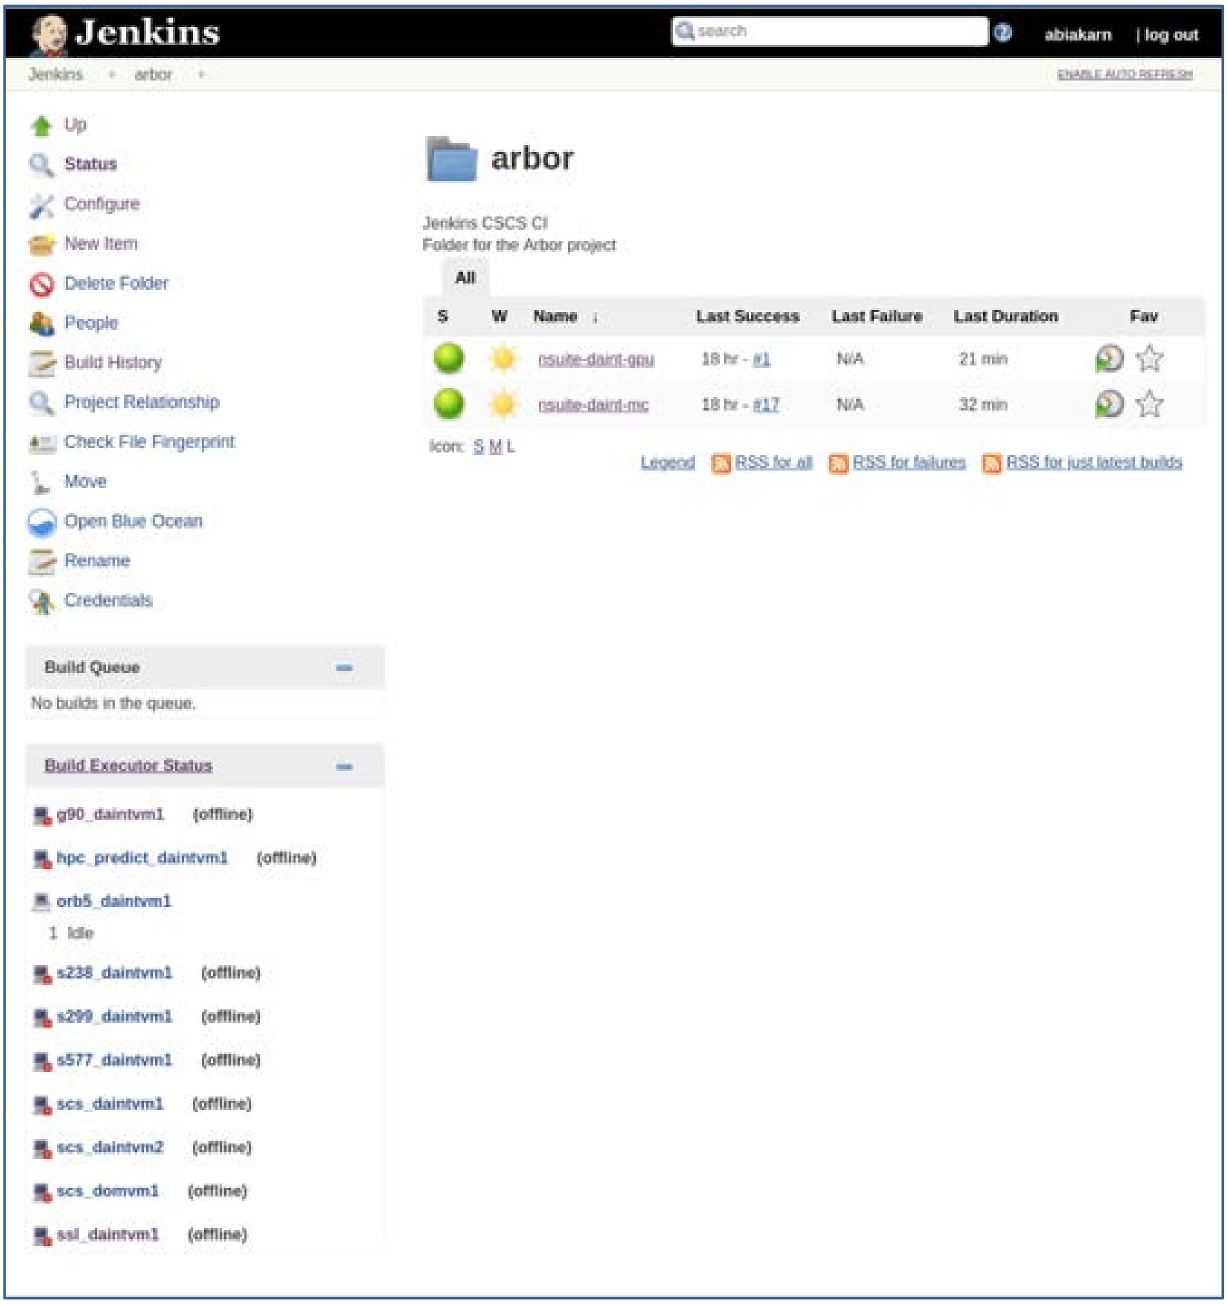

Figure 8: CSCS Jenkins CI service in use for Arbor development

7.1.6 Platform Integration Services

The UNICORE software suite received several update releases. Work has started towards the next major release (UNICORE 8), focussing on simpler deployment and easier maintenance of the software. The usage of UNICORE by other groups in the HBP was supported. For example, UNICORE is used heavily by the Brain Simulation Platform (SP6) to access HPC resources at JUELICH-JSC and ETHZ-CSCS from the HBP Collaboratory.

7.1.7 Performance optimisation

The work on performance optimisation started by establishing the methodology on how to select the codes to be optimised, and how to do a proper analysis and optimisation of them. We used PoP CoE[14]

[14] Performance Optimisation and Productivity – A Centre of Excellence in HPC (https://pop-coe.eu/)

as a source of inspiration to establish the mentioned methodology and adapted it to the particular case of neuroscience. The next step was to promote the service in the HBP, using newsletters, mailing lists, contacting specific groups and with a poster at the HBP Summit 2018. As of M12 we have made two detailed analyses of NEST and CoreNEURON codes, providing hints to the code programmers on how to orient their coding to improve performance in the future, as well as establishing good practises for the end users running the code to provide inputs in an efficient manner for achieving good performance. Ongoing discussions with other groups will provide us more applications to be analysed in the second year of SGA2.

7.1.8 User support and documentation

The HPAC user support focused during the first year of SGA2 on providing the first-level support for the HPAC infrastructure. CEA, which is a new HPAC site since the beginning of SGA2, has been included in the ticketing system used as the HPAC helpdesk. Furthermore, the new teams providing HLST support within HBP got access to the HPAC ticketing system and dedicated HLST queues have been created, so that incoming tickets can be assigned to this team, which are related to the HPAC infrastructure to be dealt with by the High-Level Support Team. The documentation of the HPAC Platform is collected in the Guidebook[15], which is updated on a regular basis.

7.1.9 Education and training

The work on education and training started by identifying the potential topics of interest and the lecturers to design the first HPAC Platform Training Event. Then, a survey was conducted to adjust the agenda to meet user needs, and work started on planning the event that took place on 11-12 December 2018 in Barcelona. In total, there were 17 attendees from eight different institutions, several of them external to HBP. During the training sessions, the attendees were introduced to the tools and services that the HPAC Platform offers, including how to get access to compute and storage resources, as well as the PRACE training. It was also explained how to use the HPAC supercomputers from the Collaboratory and how to transfer large data between sites. The simulators NEST and Arbor were presented one each with hands-on sessions and there was a last talk was about visualisation tools.

7.2 Validation and Impact

Thanks to the early access resources available at ETHZ-CSCS as part of the Fenix infrastructure, a total of eight projects have been granted compute and storage allocations already. These projects are reporting good results and some are in the process of re-applying for resources for the next twelve-month allocation period. In addition to this, a use-case driven approach is being followed whereby we use the newly gathered use case requirements coming from ICEI and WP7.1 (see Section 3.1.2) to drive the prioritisation of feature implementation and functionality improvements. As reported above these are starting to drive the work being coordinated in the weekly Joint Platform meetings. Finally, an ICEI Demonstrator is being implemented to ensure a representative end-to-end user-defined workflow that can be realised on the HPAC and Fenix infrastructure.

7.2.1 Actual Use of Output(s) / Exploitation

The HPAC Services are provided on top of the ICEI/Fenix infrastructure services, the latter providing the compute, storage and network resources needed to allow projects to produce meaningful outcomes. Hence there is a synergy between the HPAC services (and the associated human resources

[15] https://hbp-hpc-platform.fz-juelich.de/

assigned to maintain and augment them) and these ICEI/Fenix infrastructure services. Currently, a total of eight HBP projects are successfully consuming the ICEI/Fenix resources at ETHZ-CSCS.

7.2.2 Potential Use of Output(s)

The HPAC Platform has unrealised exploitation potential. There is room for considerable growth in the number of projects and users of the ICEI/Fenix resources complemented by HPAC Platform services. Furthermore, as the ICEI implementation plan is realised, yet more resources will become available as the HPAC/ICEI sites procure their systems and bring them online. Some aspects of the infrastructure have reached a high maturity level (e.g. the production HPC systems), while others are still evolving (e.g. infrastructure services interfaces for use by platforms), which leads to yet more unrealised future potential.

7.2.3 Publications

No publications yet.

7.2.4 Measures to Increase Impact of Output(s): Dissemination

More training and dissemination events are needed and the next one is scheduled for July 2019, in collaboration with the HBP Education Programme[16], and another one is planned for end of 2019. There was a noticeable impact of the ICEI/Fenix session at the HBP Summit in 2018; such efforts need to be continued.

8. Conclusion and Outlook

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

The architecture specification for the HPAC Platform and Fenix will be finalised soon. The analysis of science cases and use cases will be continued in collaboration with the respective owners, and the identified requirements will be documented. The benchmark suite will be matured into a shareable version as a collaborative effort with the ICEI project, which will be disseminated among relevant stakeholders.

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 has begun and first results are described in Deliverable D7.2.1. These developments will continue, and be validated and tested with HBP users. Work on a Central Data Location Service supporting basic functionalities for data discovery and access has started, which will be continued and released in the second year of SGA2.

Foundation work on third-factor synaptic plasticity in NEST will be the basis for the implementation of more complex plasticity rules in the second half of SGA2. The conceptual work on user-level documentation for NEST is the basis for the creation of user-level documentation for SGA3. As planned, except for NESTML advancements, little has been done in SGA2 to improve the code

[16] https://www.humanbrainproject.eu/en/education/participatecollaborate/curriculum/workshops/3rd-curriculum-workshop-ict-hpc/

sustainability of NEST; this needs to be done in SGA3. The advancements of NESTML in SGA2 enable an improvement of code sustainability of the NEST source code with abstracted model representations, which will be implemented in SGA3. The efforts related to simulator development will be continued in SGA2 and beyond.

Arbor has seen two numbered version releases, and a benchmark and validation suite has been released. Benchmarks have shown that Arbor reduces time to solution and memory overheads by an order of magnitude over previously existing simulation engines, and can run on a wider variety of HPC architectures. This simulator has a rich API that can support both model construction and model simulation with low memory and computational overheads. This work will be continued in SGA2 and this API will be used to implement support for open community standard model descriptions like SONATA, and for coupling with other simulators. This support is considered to be used in SGA3 for integrating Arbor into vertical workflows, to deliver state-of-the-art simulation capabilities as part of the HPAC Platform.

The design of the domain-specific language for high level representation of TVB models is ready and sets up the basis for the work to be completed in the second part of SGA2. Based on this, work on Bayesian inference models is planned for SGA3. The improvements with regards to the exploration capacity and performance provided by HPC kernels of TVB will allow faster and computationally more efficient model personalisation for applications inside and outside of the HBP. This will allow exploration of a larger part of the parameter space for a more accurate fitting of models to experimental data, and hence a better understanding of the model sensitivity to parameters. With this, the analysis of larger datasets and simulations with higher output prediction accuracy will be possible, which can be taken to clinical setups.

The work on NeuGen builds the basis for a user-friendly benchmark providing geometries for simulators. The coupling of multiple simulators bridges microscopic to macroscopic scales and provides the basis for larger computations. Multi-scale simulations will play an even more important role in SGA3; recent efforts in SGA2 have laid the foundations for this.

In SGA2, simulator and visualisation development in SP7 are more tightly connected than in previous phases, facilitated by merging previously separate Work Packages for these topics into one. The *in situ* pipeline has been continuously extended to support NEST and Arbor. Support for TVB is currently under development. With regards to visualisation back end development, a family of libraries has been implemented and released. It supports application-agnostic streaming capabilities to make standalone desktop visualisation applications accessible through the web. Next steps are the deployment and use on HPC as well as the integration in specific use cases.

Also, the HPAC Platform itself has seen major changes in the last year. Most notably, the migration towards the Fenix infrastructure has started and the integration of the TGCC of CEA as fifth supercomputing centre in the HBP (and Fenix) has started. Work is mostly on target, apart from some technical challenges to be solved with regards to the integration of CEA. However, these delays will not have an impact on other parts of the HBP.

The ICEI resources at ETHZ-CSCS came on-line according to the agreed schedule and some HBP users very quickly started to make use of them, supported by the HPAC and ICEI teams. There is more capacity available. However, more dissemination and outreach are needed and planned as a joint HPAC-ICEI effort to increase the uptake of the available resources. Important for that will be the finalisation of the architecture specification document.

To increase the number of users from inside and outside HBP, the HPAC Platform started to increase the dissemination and training activities, often as joint efforts together with the ICEI project. The 1st HPAC Platform Training took already place in December 2018 in Barcelona. An HBP Education Workshop “High Performance Computing for Neuroscience” is planned for July 2019.

By the end of SGA2, ICEI resources will be available at all five sites. With that, the HPAC Platform and Fenix will be the basis for the HBP Research Infrastructure EBRAINS in SGA3, also providing lower-level services, simulation engines optimised for modern and future HPC architectures, and visualisation software integrated with Fenix and HPAC.

Annex A: Additional information about releases and major updates

In the following additional information about releases and major updates are summarised in separate tables, one per Key Result.

Table 6: Overview of releases and major updates related to Key Result KR7.1

ID	Component Name	Type	Contact	Info on releases and major updates
C2668	HPAC/Fenix architecture specification	Service	JUELICH: Pleiter, Dirk	Release delayed until M13
C2669	HPAC/Fenix SGA2 architecture document v01	Report	JUELICH: Pleiter, Dirk	Release delayed until M13
C2670	HPAC/Fenix SGA2 architecture document v02	Report	JUELICH: Pleiter, Dirk	Planned for M18
C2671	SP7 use case and requirements documentation	Service	JUELICH: Klijn, Wouter	Available as living, internal documents
C2672	HPAC Platform validation (SGA2)	Service	CINECA: Testi, Debora	Internal, confidential report released in M7
C2673	HPAC SGA2 Platform validation report	Report	CINECA: Testi, Debora	Planned for M24
C2674	SP7 co-design knowledge management (SGA2)	Service	BSC: Bartolome, Javier	Internal report released in M7: https://emdesk.humanbrainproject.eu/shared/5bc98b3505114-b95af40196268becc8da5e1ce1823646

Table 7: Overview of releases and major updates related to Key Result KR7.2

ID	Component Name	Type	Contact	Info on releases and major updates
C416	Transfer services	Service	CINECA: Fiameni, Giuseppe	First version released in M11 Details are described in an internal report (related to MS7.2.1). Contact hpac-support@humanbrainproject.eu to a copy of this document.
C2823	APIs to access key technologies for interactive data analysis	Software	JUELICH: Hater, Thorsten	See Deliverable D7.2.1 for details
C2824	SLURM plug-in for the co-allocation of compute and data resources	Software	TUDA: Wolf, Felix	See Deliverable D7.2.1 for details
C2825	User-managed virtual cluster on an HPC system through SLURM	Software	CEA: Wiber, Gilles	In progress. PCOCC is available here: https://github.com/cea-hpc/pcocc
C2826	Central Data Location Service	Service	CINECA: Fiameni, Giuseppe	In progress

C2827	Report on data analytics framework performance	Report	CINECA: Rorro, Marco	In progress

Table 8: Overview of releases and major updates related to Key Result KR7.3

ID	Component Name	Type	Contact	Info on releases and major updates
C209	NEST - The Neural Simulation Tool	Software	JUELICH: Diesmann, Markus	NEST release 2.16.0 http://www.nest-simulator.org/download/#releases
C510	Continuous dynamics code in NEST	Software	JUELICH: Diesmann, Markus	NEST functionality has been considerably enhanced, including the adaptation of the rate model framework for novel use cases (released in NEST 2.16.0, see C209)
C661	NEST Support for Providers	Service	NMBU: Plesser, Hans Ekkehard	Extensive notes were produced for each release and communication with the software maintainers of HPAC and NRP continued
C662	NEST Requirements Management	Service	NMBU: Plesser, Hans Ekkehard	The software framework for synaptic plasticity was extended to prepare for the representation of third factor plasticity and a corresponding pull request including a simple example (Clopath plasticity) made against and integrated in NEST. https://github.com/nest/nest-simulator
C2678	Prototype implementation of unified connectivity generation scheme for very large networks	Software	NMBU: Plesser, Hans Ekkehard	Unified representation of network structure with and without spatial structure and significantly improved threading performance for connecting complex networks implemented as prototype; user interface for unified connectivity generation scheme drafted.
C2679	Prototype NEST code with abstracted synapse model representations	Software	JUELICH: Morrison, Abigail	The high-level modelling language, NESTML, has been advanced to permit synapse models to be expressed as well as neuron models. https://github.com/nest/nestml
C2689	Publication on Arbor methods and implementation	Report	ETHZ: Cumming, Benjamin	Conference presentation at PDP 2019, in Pavia Italy, on Feb 15 2019. The paper will be published in the associated conference proceedings: "Arbor - A Morphologically-Detailed Neural Network Simulation Library for Contemporary High-Performance Computing Architectures".
C2690	HPC multi-compartment benchmark and validation suite	Software	ETHZ: Cumming, Benjamin	Released with Deliverable D7.3.1 (D44.1, D83) "Multi-compartment benchmark and validation suites on HPC systems" (submitted on 6 March 2019).

				Available online at www.github.com/arbor-sim/nsuite
C2691	Arbor library and frameworks (SGA2 release)	Software	JUELICH: Peyser, Alexander	Versions 0.1 and 0.2 released in M7 and M12, respectively, available on GitHub: www.github.com/arbor-sim/arbor Major features implemented: • Gap junctions • Back end optimisations (Intel SIMD code generation, fast parallel GPU matrix solver) • Management of hardware resources (GPUs, MPI, threads)
C2692	NeuGen: Interactive neuronal network geometry generation tool	Software	UFRA: Wittum, Gabriel	NeuGen has been coupled to VRL Studio in order to provide a convenient user-interface. NeuGen: https://github.com/NeuroBox3D/NeuGen Latest release: 30 Jul 2016 VRL: https://github.com/VRL-Studio/VRL Latest release: 31 October 2018 VRL Studio: https://github.com/VRL-Studio/VRL-Studio Latest release: 30. October 2018
C2693	Simulations with NEST – NEURON coupling using MUSIC	Report	UFRA: Wittum, Gabriel	Simple simulation with NEST and NEURON and UG4 has been established.
C2694	Development of a model with full spatial resolution for the extracellular potential activity using ug4	Model	UFRA: Wittum, Gabriel	A physically realistic fully resolved model (3D resolution) based on Maxwell's laws and ohmic currents has been derived and implemented. Under testing. https://github.com/UG4/ugcore Latest release: 3. April 2018, Pre-release: 4. January 2019
C2695	Simulations showing extracellular potential activity as result of an active network which couple the model with full spatial resolution (implemented in ug4) with MUSIC	Software	UFRA: Wittum, Gabriel	Simple simulation with NEST and NEURON and UG4 has been established.
C2696	Prototype implementation of framework and numerics for phenomenological neuron model with compartmentalized third factor plasticity	Software	JUELICH: Diesmann, Markus	The software framework for synaptic plasticity was extended to prepare for the representation of third factor plasticity and a corresponding pull request including a simple example (Clopath plasticity) made against and integrated in NEST.
C2789	NEST user-level documentation	Software	NMBU: Plesser, Hans Ekkehard	On 21-22 November 2018 the workshop "Are we building the right user-level documentation" was held (T7.4.5) at EITN together with the developers of neuromorphic hardware (SP9).

				A summary of the report can be provided on demand (hpac-support@humanbrainproject.eu).
C3030	HPC-for-TVB software package	Software	JUELICH: Diaz, Sandra	The design for the DSL has been tested together with tools to visualise the correct translation of the models into executable kernels. No formal release is available yet. The current status of the code can be found in the GitHub repository (https://github.com/the-virtual-brain/tvb-hpc/).

Table 9: Overview of releases and major updates related to Key Result KR7.4

ID	Component Name	Type	Contact	Info on releases and major updates
C2680	HPAC integration of in situ pipeline	Software	RWTH: Weyers, Benjamin	https://devhub.vr.rwth-aachen.de/VR-Group/contra https://devhub.vr.rwth-aachen.de/VR-Group/nesci First release: 18.07 (July 2018) Support for streaming NEST data via shared memory including python bindings. Major updates since first release: Added support for Arbor and transport via ZeroMQ as well as simplifying the infrastructure work to integrate the pipeline into external projects.
C2681	Development, extension and deployment of visualisation tools for HPAC	Software	URJC: Pastor, Luis	Deployment in CINECA (Galileo) ViSimpl: Visual discrimination of groups of neurons; neuron picking; visual and render enhancements NeuroTessMesh: different visual codification for different parts of a morphology; improvements for dealing with occlusions; information on neuron activity; OpenMP-based optimisation
C2682	Remote high-fidelity visualisation and interaction middleware	Software	URJC: Pastor, Luis	ReMo (https://github.com/gmrvvis/ReMo): low-level library for managing video streams and connecting video generators with video consumers. ReMoToo (https://github.com/gmrvvis/ReMoToo): service running in the visualisation nodes which relies in ReMo+webstreamer capabilities for capturing the desktop and stream it to the client. ReMoLON (https://github.com/gmrvvis/ReMoLON): composed of two services: the front-end and the “instance manager”.

C2697	Flexible in situ pipeline coupling NEST to visual analysis tools	Software	RWTH: Weyers, Benjamin	No official release yet. In the current state, the streaming module offers a simple way to connect the in situ pipeline to the NEST simulator without any source code modification. https://devhub.vr.rwth-aachen.de/VR-Group/nest-streaming-module
C2698	Evolution of visual analysis tools towards the in-situ pipeline	Software	URJC: Pastor, Luis	ConGen: define populations of cells; establish connections among them; creation of inputs and connection to populations
C2707	Semantic Visualisation Software Catalogue	Software	RWTH: Weyers, Benjamin	First version of architecture and catalogue content available: https://github.com/bweyers/HBPVisCatalogue

Table 10: Overview of releases and major updates related to Key Result KR7.5

ID	Component Name	Type	Contact	Info on releases and major updates
C263	UNICORE	Software	JUELICH: Schuller, Bernd	Maintenance releases https://www.unicore.eu/download/
C329	SP7 Federated HPAC Data and Computing Services	Service	JUELICH: Lippert, Thomas	Business as usual operations
C337	HPAC Platform Guidebook for SGA2	Report	JUELICH: Lührs, Anna	Regular updates of the public website: https://hbp-hpc-platform.fz-juelich.de/
C409	Archive data repositories	Service	CINECA: Mucci, Roberto	Business as usual operations
C410	Active data repositories	Service	ETHZ: McMurtrie, Colin John	Business as usual operations
C416	Transfer services	Service	CINECA: Fiameni, Giuseppe	Testing and development
C417	Data federation and management protocols	Service	CINECA: Fiameni, Giuseppe	Refinement of specifications
C571	Accounting service	Service	JUELICH: Krause, Dorian	Business as usual operations
C1093	CSCS Reporting and Accounting Service	Service	ETHZ: McMurtrie, Colin John	Business as usual operations
C1094	JSC Reporting and Accounting Service	Service	JUELICH: Lippert, Thomas	Business as usual operations
C1096	CINECA Reporting and Accounting Service	Service	CINECA: Fiameni, Giuseppe	Business as usual operations
C1097	BSC Reporting and Accounting Service	Service	BSC: Bartolome, Javier	Business as usual operations

C1112	HPAC Data Services	Service	ETHZ: McMurtrie, Colin John	Business as usual operations and refinement
C1114	HPAC Ticket System	Software	BSC: Vicente, David	Business as usual operations
C1129	CINECA Archive Repository Service	Service	CINECA: Fiameni, Giuseppe	Business as usual operations
C2709	Intermediate code optimisation report	Report	BSC: Sirvent, Raül	Delivered
C2710	Final code optimisation report	Report	BSC: Sirvent, Raül	Planned for M24
C2828	CEA Computing Service	Service	CEA: Lafoucriere, Jacques-Charles	Integration in progress
C2831	CEA Monitoring Service	Service	CEA: Lafoucriere, Jacques-Charles	Integration in progress
C2832	CEA AAI Service	Service	CEA: Lafoucriere, Jacques-Charles	Integration in progress
C2834	CEA Reporting and Accounting Service	Service	CEA: Lafoucriere, Jacques-Charles	Integration in progress
C2837	CEA Data Service	Service	CEA: Lafoucriere, Jacques-Charles	Integration in progress
C2838	CEA Active Repository Service	Service	CEA: Lafoucriere, Jacques-Charles	Integration in progress
C2839	BSC Active Repository Service	Service	BSC: Bartolome, Javier	Integration in progress
C2840	CEA Archive Repository Service	Service	CEA: Lafoucriere, Jacques-Charles	Integration in progress
C2841	BSC Archive Repository Service	Service	BSC: Bartolome, Javier	Business as usual operations
C2842	CEA Upload/Download Service	Service	CEA: Lafoucriere, Jacques-Charles	Integration in progress
C2843	BSC Upload/Download Service	Service	BSC: Bartolome, Javier	Optimisation
C2849	CEA LDAP	Software	CEA: Lafoucriere, Jacques-Charles	Integration in progress
C2855	CEA Resource Usage Tools	Software	CEA: Lafoucriere, Jacques-Charles	Integration in progress
C2860	RWTH visualisation cluster	Hardware	RWTH: Weyers, Benjamin	Operational
C2861	RWTH storage system	Hardware	RWTH: Weyers, Benjamin	Operational
C2867	HPAC Platform training programme	Service	BSC: Sancho, Maria Ribera	Underway

				https://hbp-hpc-platform.fz-juelich.de/?page_id=583 1st HPAC Platform Training: https://www.bsc.es/education/training/other-training/1st-hpac-platform-training-event “HPC for Neuroscience” Education Workshop: https://www.humanbrainproject.eu/en/education/participatecollaborate/curriculum/workshops/3rd-curriculum-workshop-ict-hpc/
C2870	HPAC Platform operational status report	Report	ETHZ: McMurtrie, Colin John	In progress
C2871	HPAC Platform resource allocation policy summary	Report	CEA: Lafoucriere, Jacques-Charles	In progress
C2872	HPAC Platform security policy review	Report	ETHZ: McMurtrie, Colin John	In progress
C2873	HPAC Platform Review of Data Management Plan and Data Policy Manual documents	Report	ETHZ: McMurtrie, Colin John	In progress
C2875	HPAC Platform lightweight virtualisation environment definition report	Report	ETHZ: McMurtrie, Colin John	In progress