



# VirtualBrainCloud

Personalized Recommendations for  
Neurodegenerative Disease



UNICORE

FENIX  
RESEARCH INFRASTRUCTURE



[www.VirtualBrainCloud-2020.eu](http://www.VirtualBrainCloud-2020.eu)

## Public deliverable report

D7.8: Gateway to PRACE documentation and guidelines

Date	January 2020
Authors	Michael Tarnawa, Sandra Diaz (Forschungszentrum Jülich FZJ) © VirtualBrainCloud consortium
Dissemination level	<b>public</b>
Website	<a href="http://www.VirtualBrainCloud-2020.eu">www.VirtualBrainCloud-2020.eu</a>



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826421



## Table of content

1. Introduction.....	4
1. Requesting Access .....	4
1.1.1. PRACE .....	4
1.1. ICEI/Fenix.....	8
2. Usage of HPC Systems .....	9
2.1. Standard Connection.....	10
2.2. UNICORE.....	11
3. Conclusion .....	13



## List of Abbreviations

Abbreviation	Explanation
BSC	Barcelona Supercomputing Center
CSCS	Swiss National Supercomputing Centre
ETH Zurich	Eidgenössische Technische Hochschule Zürich
FTP	File Transfer Protocol
GCS	Gauss Centre for Supercomputing e. V.
GENCI	Grand équipement national de calcul intensif
GPFS	IBM Spectrum Scale (f/k/a General Parallel File System)
GUI	Graphical User Interface
HBP	Human Brain Project
HPC	High Performance Computing
ICEI	Interactive Computing E-Infrastructure
JSC	Jülich Supercomputing Centre
PI	Principal Investigator
PMO	ICEI Project Management Office
PRACE	Partnership for Advanced Computing in Europe
R&D	Research and development
REST	Representational state transfer
SCP	Secure Copy
Slurm	Simple Linux Utility for Resource Management
SSH	Secure Shell
TVB-Cloud	VirtualBrainCloud
UNICORE	Uniform Interface to Computing Resources
URL	Uniform Resource Locator
VM	Virtual machine



## 1. Introduction

The objective of the Virtual Brain Cloud project is to develop TVB-Cloud, a cloud-based platform for biomedical research and clinical decision-making that will help to improve early patient-specific diagnosis and treatment of neurodegenerative diseases. This cloud-based environment will rely on big data and high-performance computing (HPC) to achieve its goals. For more information about Virtual Brain Cloud, please visit the official project website<sup>1</sup>.

In science and industry, compute intensive processes are typically executed on computer clusters. Computer clusters have relatively high acquisition and maintenance costs for a single institution. This has led to the emergence of High Performance Computing centres that provide distributed access to supercomputer systems for several users. The TVB-Cloud infrastructure will rely on HPC services for compute intense calculations. It will offer its users the option to send their pipelines to available supercomputers. For this reason, a gateway to the systems needs to be established first, following the standard regulations by the HPC centres.

This document describes the general formalities needed for accessing HPC systems in Europe, namely, the application process in PRACE (Partnership for Advanced Computing in Europe) and ICEI/Fenix (Interactive Computing E-Infrastructure). Thereafter, the connection to the HPC systems is described from the user perspective, introducing the federated middleware UNICORE (Uniform Interface to Computing Resources), which provides, apart from Secure Shell, a method for creating, submitting and monitoring compute jobs on the various computing systems of PRACE.

## 1. Requesting Access

As the demand from users is greater than the total number of resources at the HPC sites, the access to the offered systems needs to be regulated, for instance, by restricting the amount of computing time per user. For projects with international users like the Virtual Brain Cloud, this is realised in a peer-review process organised by PRACE, which was established to facilitate the access to the systems and the application process for computing time on the different HPC systems in Europe. In addition, five of the HPC centres collaborate in the ICEI project to build the Fenix Infrastructure. Neuroscientists can apply through the mechanisms provided by PRACE and ICEI/Fenix for access to the computer systems.

### 1.1.1. PRACE

PRACE<sup>2</sup> is an association of 26 European countries with its head office in Brussels and offers computing and data management resources and services at various HPC sites. The sites are divided in two tiers. Tier-0 contains larger systems from the five hosting members<sup>3</sup>: BSC in Spain, CINECA in Italy, ETH Zurich/CSCS in Switzerland, GCS in Germany and GENCI in France. Tier-1 consists of smaller systems

---

<sup>1</sup> <https://virtualbraincloud-2020.eu/>

<sup>2</sup> <http://www.prace-ri.eu/>

<sup>3</sup> <http://www.prace-ri.eu/prace-resources/>



across Europe. These HPC-systems can in general be accessed only by applying for compute time with a proposal in the peer-review process.

Scientists and researchers from academia and industry from around the world can choose between two access schemes to get one of the PRACE resources: Preparatory Access and Project Access. The former allows a short-term access to the resources for a project. During this time, programming codes can be tested for scalability on the supercomputers. Preparatory Access is categorized in types A to D, having different purposes, support, and run times<sup>4</sup>:

- **Type A:** This scheme is intended to produce scalability plots of the performance of the codes on PRACE HPC systems, as well as other parameters that may be relevant to apply for PRACE calls for Project Access. The maximum duration of Type A projects is 2 months.
- **Type B:** The objective of this scheme is to undertake code development and optimisation. Applicants will need to describe the proposed work plan in detail, including the human resources and expertise available to implement the project. The maximum duration of Type B projects is 6 months.
- **Type C:** In this scheme, PRACE experts are requested to provide the necessary support to undertake adaptations (development and optimisation) to the codes of PRACE users. The maximum support that can be requested is the equivalent to 6 person-months of effort. The maximum duration of Type C projects is 6 months.
- **Type D:** This scheme allows PRACE users to start a code adaptation and optimisation process on a PRACE Tier-1 system. PRACE experts help in the system selection process. In addition to Tier-1 computing time, the PRACE user will also receive Tier-0 computing time towards the end of the project (inform of the Type A scheme) to test the scalability improvements. The work is supported by PRACE experts. The maximum support that can be requested is the equivalent to 6 person-months of effort. The maximum duration of Type D projects is 12 months.

Type A and Type B proposals can be submitted anytime and the access will be available within two weeks. Proposals for Type C and Type D, however, are evaluated and granted quarterly only.

PRACE offers an online form for submissions, which has to meet certain general requirements. The project must have a civilian purpose and the project leader shall have an employment contract with its institution valid for at least three months after the end of the allocation period. Additionally, if the researcher is from a private company, the head office or substantial R&D have to be located in Europe and they have to use the resources only for open R&D research purposes. Technical evaluations are performed by recognised technical experts from the relevant PRACE Hosting Members on the applications. For Type C and Type D proposals, software experts will be also involved in the review process.

Proposals for Preparatory Access Types B and C must provide a detailed description of the issues hindering the scalability of the code to higher computing performance, so as to facilitate the technical

---

<sup>4</sup> <http://www.prace-ri.eu/prace-preparatory-access/>



assessment. In addition to this, a clear plan for code development must be proposed and the available human resources and expertise (available and/or required) to execute this planning must be quantified. Proposals for Preparatory Access Type D must provide information about the system(s) used so far and a short description of scalability issues already discovered (if known). In addition, similar to Type B and C, a clear work plan for code development must be proposed and the available human resources and expertise (available and/or required) must be specified.

The awarding of the Preparatory Access proposals is limited by the amount of resources available. For those proposals with a positive technical evaluation, awarding will be based on a first-come, first-served basis, according to the final submission date. For Type C and D, there are also limitations regarding the availability of the support from PRACE experts, in terms of staff hours, and the type of expertise required.

The Project Access<sup>5</sup> is recommended for long-time productive runs on already tested code. It offers 12-month, 24-month, and 36-month solutions to choose from. Similar to Preparatory Access, only projects with open R&D research purposes are considered. A proposal is comprised of two forms, the Online Application form and the Project Scope and Plan, and must be submitted before the call deadline, which is set twice a year. It is also possible to apply for ICEI resources.

The PRACE peer review process is based on the following principles<sup>6</sup>:

- **Transparency:** The peer review process is transparent and clear to all stakeholders in PRACE including funding agencies of all member countries and users from research institutions and industry.
- **Fairness:** Proposals are evaluated on merit and potential high impact on European and international science and economy.
- **No parallel assessment:** The PRACE peer review process builds on the experiences and best practices of national and international institutions and constitutes a centralised peer review exercise recognised by all PRACE partner countries and scientific communities.
- **Reviews** are done by experts in the scientific field of the proposal, with no declared conflict of interest, based on criteria published in the PRACE Calls for Proposals and with a periodic reshuffling of reviewer's appointments.
- **Confidentiality:** Proposals will be treated with the needed confidentiality by PRACE staff and reviewers. The identities of the peer reviewers shall not be disclosed.
- **Right to reply** to technical and scientific evaluations.

The assessment procedure encompasses four phases<sup>7</sup> (cf. figure 1):

- **Administrative check (phase I).** Proposals not complying with PRACE eligibility criteria will be rejected at this stage and will not continue to the next phase.

---

<sup>5</sup> <http://www.prace-ri.eu/prace-project-access/>

<sup>6</sup> <http://www.prace-ri.eu/peer-review/>

<sup>7</sup> [http://www.prace-ri.eu/IMG/pdf/Terms\\_of\\_Reference\\_Call20-1.pdf](http://www.prace-ri.eu/IMG/pdf/Terms_of_Reference_Call20-1.pdf)



- **Technical Assessment (phase II).** Proposals will be technically reviewed by technical experts of PRACE Hosting Member sites. During this phase, applicants may be contacted by technical experts in case of questions or concerns raised during the review.
- **Scientific Assessment (phase III).** Proposals will be peer reviewed by recognised independent scientific experts. Individual reports will be made available to applicants. Applicants are strongly recommended to use the opportunity to comment on these assessments during the right to reply period.
- **Access Committee Assessment (phase IV).** The reviewers' reports and the applicants' responses will be analysed by the PRACE Access Committee who will produce the final ranking list. The Access Committee may agree on a scientific quality cut-off threshold. Proposals ranked under this threshold will not be awarded, even if there are resources available on the systems.

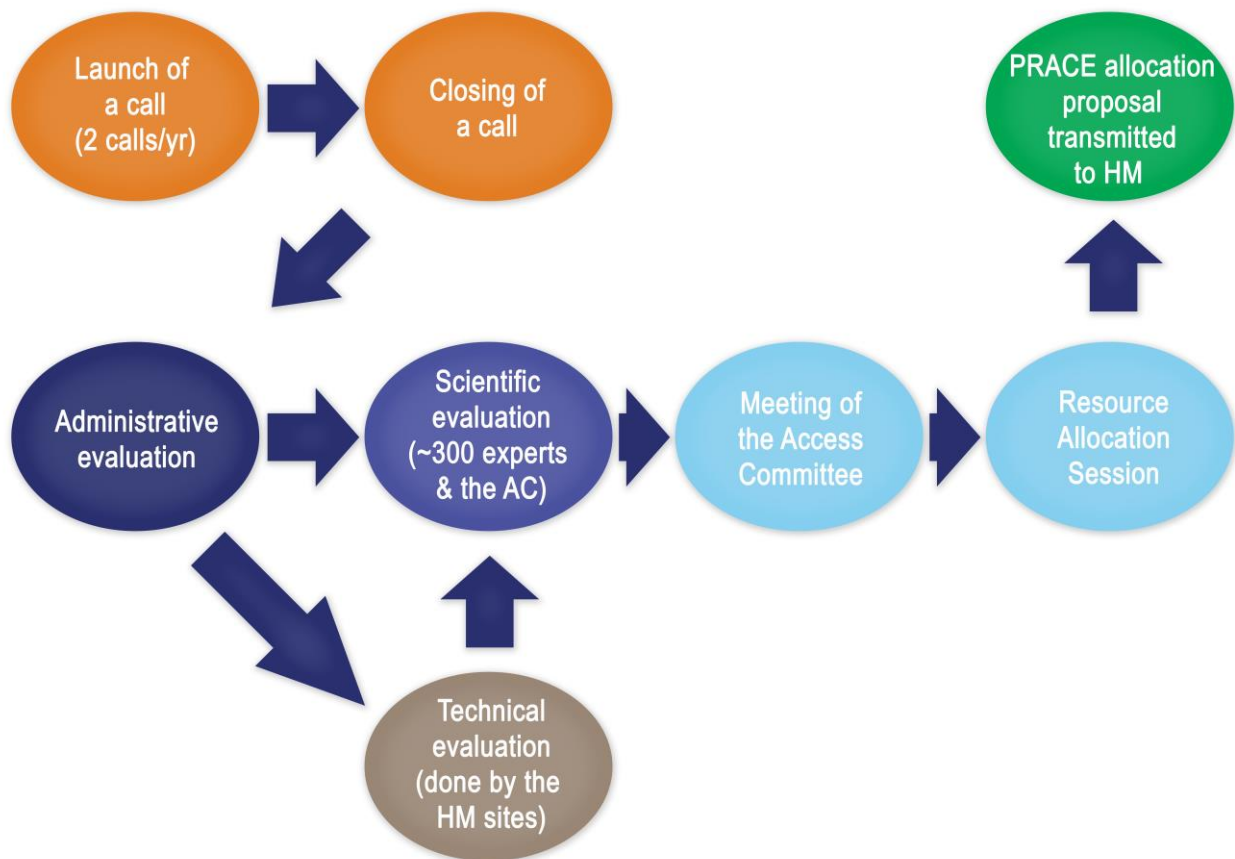


Figure 1: PRACE reviewing process on a proposal (<http://www.prace-ri.eu/prace-peer-review/>)

For the technical and scientific assessment, a set of criteria is specified and should be fully addressed in the application. Reviewers will assess proposals against these criteria. A proposal must convincingly address the following criteria in the technical part:

- The **need** to use PRACE resources.
- **Software** availability on the requested resource. The codes necessary for the project must already be available on the system requested, or in case of codes developed by the applicants and Multi-year and Project Access proposals, they must have been sufficiently tested for efficiency, high scalability, and suitability. For Multi-year and Project Access Proposals proof of successful tests must be submitted together with the proposal.



- **Feasibility** of the requested resource. The requested system(s) must be suitable for the proposed project. The technical assessment may redirect projects to a more appropriate system.

The proposals must address the following scientific criteria:

- **Scientific excellence.** The proposed research must demonstrate scientific excellence and a potential for high European and international impact.
- **Novelty and transformative qualities.** Proposals should be novel, develop an important scientific topic of major relevance to European research, describe possible transformative aspects, and expected advances.
- **Relevance to the call.** Proposal needs to address how the proposed project is addressing the scope of the call if a specific scope is stated in the call.
- **Methodology.** The methodology used should be described and appropriate to achieve the goals of the project.
- **Dissemination.** The planned channels and resources for dissemination and knowledge exchange should be described.
- **Management.** There must be a solid management structure which will ensure that the project will be successfully completed.

Near the end of any resource access period (i.e., Preparatory Access and Project Access), the research group must commit a final report to PRACE. Failure to submit a final report may disqualify future proposal submissions to PRACE by any member of the research group. Also, applicants must acknowledge PRACE in all publications that describe results obtained using PRACE resources.

### 1.1. ICEI/Fenix

Fenix<sup>8</sup> is a consortium of five European supercomputing centres, namely BSC (Spain), CEA (France), Cineca (Italy), CSCS (Switzerland), and JSC (Germany), that have agreed to build the Fenix Infrastructure. An initial version of the Fenix Infrastructure is created by the consortium within the ICEI (Interactive Computing E-Infrastructure for the Human Brain Project) project. The Fenix Infrastructure offers several services<sup>9</sup>:

- **Interactive Computing Services**  
Quick access to single compute servers to analyse and visualise data interactively, or to connect to running simulations, which are using the scalable compute services.
- **Scalable Computing Services**  
Massively parallel HPC systems that are suitable for highly parallel brain simulations or for high-throughput data analysis tasks.
- **Virtual Machine Services**

---

<sup>8</sup> <https://fenix-ri.eu/>

<sup>9</sup> <https://fenix-ri.eu/infrastructure/services>





Service for deploying virtual machines (VMs) in a stable and controlled environment that is, for example, suitable for deploying platform services like the HBP Collaboratory, image services or neuromorphic computing front-end services.

- **Active Data Repositories**

Site-local data repositories close to computational and/or visualization resources that are used for storing temporary replicas of data sets. In the near future, they will typically be realised using parallel file systems.

- **Archival Data Repositories**

Federated data storage, optimized for capacity, reliability and availability that is used for long-term storage of large data sets, which cannot be easily regenerated. These data stores allow the sharing of data with other researchers inside and outside of HBP.

There exist two ways to get access<sup>10</sup> to the available resources. The first one is to submit an application to PRACE. There, the option to apply for the ICEI/Fenix resources has to be chosen.

The other way is the HBP allocation mechanism, which is available to members of the Human Brain Project<sup>11</sup> (HBP) and partnering projects such as the Virtual Brain Cloud project, as well as neuroscientists outside of the HBP. Fenix provides a template<sup>12</sup> that research groups can use to prepare a proposal. Proposals can be submitted to the ICEI Project Management Office (PMO) at any time and will be technically and scientifically reviewed by experts from ICEI and EBRAINS<sup>13</sup> following the PRACE peer review principles. In the case of a positive evaluation, access will in general be granted within one month after the submission.

## 2. Usage of HPC Systems

When the project proposal has been granted, the required resources become accessible on the requested HPC system. Eventually, the project leader, also known as Principle Investigator (PI), will be notified of the completion of the proposal assessment and receives instructions on how to apply for user accounts on the HPC system. From there, PIs can manage their project themselves. At JSC, for example, this is managed on JuDoor<sup>14</sup>, where a PI can check the usage statistics and can invite other researchers to participate in the project. When an account on an HPC system has been created, the user can access the HPC system remotely. Normally, they are connected to the login nodes of the cluster and are able to store their data on a file system before executing the code on the compute nodes. Eventually, the results are transferred back to their local machine. As distinct infrastructures are used, the operation on the various computing centres differs in some details, for instance, data storage and job execution. Middleware like UNICORE unify the interface for accessing HPC systems, adjusting the differences between them, such that a user does not have to adapt to a new structure.

---

<sup>10</sup> <https://fenix-ri.eu/access>

<sup>11</sup> <https://www.humanbrainproject.eu/en/>

<sup>12</sup> [https://fenix-ri.eu/sites/default/files/public/file-uploads/hbp-icei-full-mechanism\\_resource-request-v07\\_1.docx](https://fenix-ri.eu/sites/default/files/public/file-uploads/hbp-icei-full-mechanism_resource-request-v07_1.docx)

<sup>13</sup> <https://ebrains.eu/>

<sup>14</sup> <https://judoor.fz-juelich.de/login>



## 2.1. Standard Connection

In general, computing centres offer remote access to their resources over Secure Shell (SSH). A user can connect to the login nodes of the supercomputer with their username and password. The authentication of the host must be verified once on the first login. Most systems also allow the exchange of SSH keys for an easier and more secure way for authentication. After logging in, the interactive shell of a login node is available. From there, the user has access to storage space where they can put their data and compile their code. The directory structure is specified by the HPC centres, some of them offering dedicated nodes for data transfer. The commonly used ways for transfer are secure copy (SCP), rsync, and File Transfer Protocol (FTP). The first two are executed from the command line, where SCP only copies files over SSH and rsync synchronises whole directories. For the exchange via FTP, there exist programs with graphical user interfaces (GUIs) on the client side.

A user cannot access the compute nodes for computation directly. The program execution on these nodes is handled by a batch system instead, which decides when and on which nodes the program will run. For doing that, the user has to write a job script which contains the instructions and additional metadata, such as the number of nodes and processes, the cut-off time, and the billing account. The submitted job will then be scheduled automatically, taking the provided information of the metadata into account. The output of a running job is usually written to a temporary storage, specially suited for I/O operations. After job completion, the generated data can then be transferred back to the user's local computer or stored for a longer time period in the supercomputing centre if archiving is allowed. The state of the storage and batch systems at JSC is described in the following paragraphs as an architectural example found at an HPC site.

All supercomputer systems, community clusters and the data access system at the Jülich Supercomputing Centre (JSC) are connected to a General Parallel File System (GPFS). Users have access to various GPFS file systems which are intended for different purposes<sup>15</sup>. Normally, a compute project is granted giving access to three file systems : \$HOME, \$SCRATCH, and \$PROJECT. If additional storage space is needed or data should be stored independent of compute projects, a separate data project needs to be applied for.

- **\$HOME**  
Acts as repository for the user's personal data like the SSH key. There is a separate HOME folder for each HPC system and a shared folder which points on all systems to the same directory. Data within \$HOME are backed up.
- **\$SCRATCH**  
Is bound to a compute project and acts as a temporary storage location with high I/O bandwidth. If the application needs to handle large files and I/O demands, \$SCRATCH is the right file system to place them. Data within \$SCRATCH is not backed up and daily cleanup is done.

---

<sup>15</sup> [https://www.fz-juelich.de/SharedDocs/FAQs/IAS/JSC/EN/JUST/FAQ\\_00\\_File\\_systems.html](https://www.fz-juelich.de/SharedDocs/FAQs/IAS/JSC/EN/JUST/FAQ_00_File_systems.html)



- **\$PROJECT**  
Is a data repository for a compute project. Its lifetime is bound to the project lifetime. Data are backed up.
- **\$FASTDATA**  
Belongs to a data project. This file system is bandwidth optimized (similar to \$SCRATCH), but data are persistent and internally backed up via snapshots.
- **\$DATA**  
Belongs to a data project. This file system is designed to store a huge amount of data on disk based storage. The bandwidth is moderate. The file-system internal backup is realized with the GPFS snapshot feature.
- **\$ARCHIVE**  
Is bound to a data project and acts as storage for all files not in use for a longer time. Data are migrated to tape storage. It is recommended to use tar-files with a minimum size of multiple Gigabytes and maximum of 8 TB. The background is that recalling/restoring files from tape is much more efficient using only a few large data streams than thousands of small data streams.

HPC systems in the JSC use the Slurm Workload Manager (Simple Linux Utility for Resource Management), a free open-source resource manager and batch system. Slurm is a modern, extensible batch system that is widely deployed around the world on clusters of various sizes. In Slurm a job is an allocation of selected resources for a specific amount of time. A job allocation can be requested using sbatch and salloc. Within a job multiple job steps can be executed using srun that use all or a subset of the allocated compute nodes. Job steps may execute at the same time if the resource allocation permits it. For more information on Slurm please refer to the online documentation<sup>16</sup>.

## 2.2. UNICORE

The federated middleware system UNICORE (Uniform Interface to Computing Resources) was developed to facilitate creating, submitting and monitoring compute jobs on the different PRACE computing sites and their batch systems. UNICORE provides seamless, uniform and secure access to the heterogeneous PRACE infrastructure. UNICORE servers are installed on systems in the supercomputing centres, allowing users access to the HPC-systems through graphical and command line clients. The latest version of the UNICORE rich client (URC) and UNICORE command line client (UCC) can be found on the official website<sup>17</sup>. The URC provides a graphical user interface (GUI) for job submission and data transfer.

Representational state transfer (REST) is another option for accessing the HPC systems through UNICORE. Interactions with a supercomputer can be managed by sending GET, POST, and DELETE statements to specific URLs. Providing an easier way, the Python library pyunicore wraps the REST API

---

<sup>16</sup> <https://slurm.schedmd.com/documentation.html>

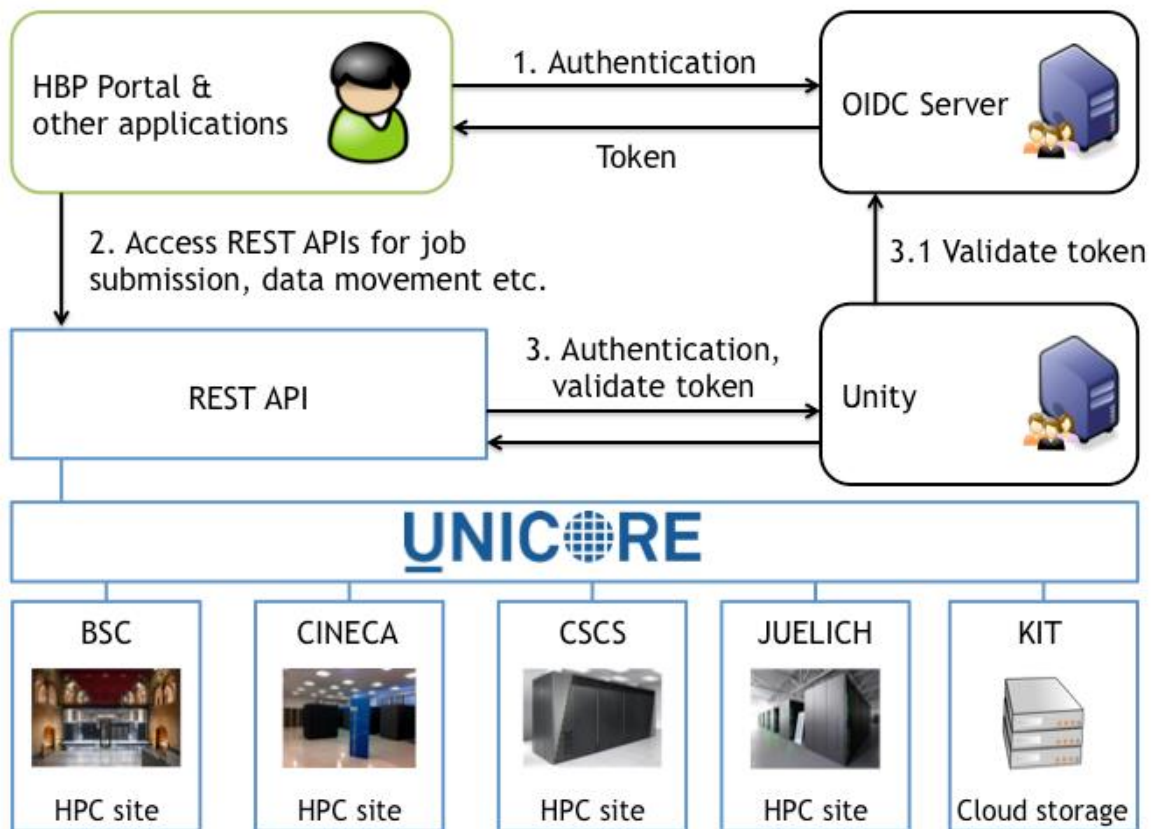
<sup>17</sup> <https://www.unicore.eu>



behind a python module. REST and pyunicore are used inside the HBP Collaboratory for job submissions to HPC centres.

In the case of HBP users, after successfully applying for compute time at the desired HPC centre, the user accounts are automatically linked to the user’s HBP OIDC account (i.e. to the Collaboratory account) so that the UNICORE portal and the UNICORE REST API can be used seamlessly from the HBP Collaboratory. Further information on the usage of the REST API can be found in the Documentation<sup>18</sup> and Examples<sup>19</sup> pages of the UNICORE Wiki on Sourceforge.

Figure 2 describes how UNICORE, the OIDC Server and Unity communicate in order to authenticate an HBP user. After a successful validation, the user can access the resources of the HPC Platform using the REST API. Another way to access the resources is via the UNICORE Portal, which is also integrated into the HPC Platform Collab in the Collaboratory.



**Figure 2** Unified access to the HPC Platform of the HBP via UNICORE and the REST API  
 ([https://hbp-hpc-platform.fz-juelich.de/?page\\_id=60](https://hbp-hpc-platform.fz-juelich.de/?page_id=60))

<sup>18</sup> [https://sourceforge.net/p/unicore/wiki/REST\\_API/](https://sourceforge.net/p/unicore/wiki/REST_API/)

<sup>19</sup> [https://sourceforge.net/p/unicore/wiki/REST\\_API\\_Examples](https://sourceforge.net/p/unicore/wiki/REST_API_Examples)



### 3. Conclusion

This document has described the steps necessary for establishing a gateway to HPC resources from a standard user perspective. For researchers across Europe, connections are regulated by PRACE. Research groups have the choice to apply for Preparatory Access or Project Access to one of the HPC systems according to their needs. The former is for testing purposes, while the latter is intended for production runs. All correctly submitted proposals are evaluated by technical and, in the case of Project Access, scientific experts assigned by PRACE, following the specified peer review principles.

Because of TVB-Cloud now being an HBP partnering community, we also introduced Fenix, an infrastructure collaboration of five PRACE HPC sites included in the ICEI project. In addition to access via PRACE, HBP partnering communities can also be granted access to HPC resources by the Fenix PMO.

This deliverable concluded with a look at the normal work flow when interacting with an HPC system remotely. As the details differ on every HPC site, some aspects in file storage and job scheduling specific for JSC were highlighted. Alternatively, UNICORE was presented as a middleware system, installed on most PRACE resources, which allows access to different compute clusters with the same commands inside a GUI or the command line.

