Final release of the Collaboratory (D6.5 – SGA3)

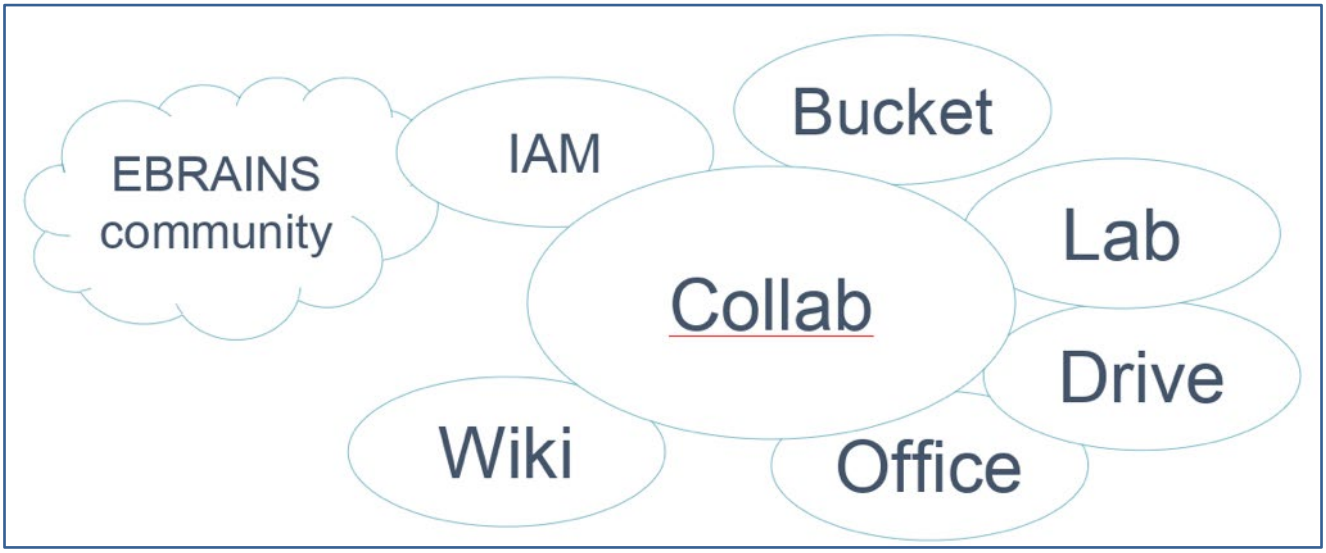

Figure 1: The Collaboratory group of services shown and how they all work together

An easy overview of everything that the Collaboratory has created and what they offer.

Project Number:	945539	Project Title:	HBP SGA3

Document Title:	D6.5 Final release of the Collaboratory
Document Filename:	D6.5 (D61) SGA3 M36 SUBMITTED 230316.docx
Deliverable Number:	SGA3 D6.5 (D61)
Deliverable Type:	Other
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA3 M36 / 31 MAR 2023
Actual Delivery Date:	SGA3 M36 / 16 MAR 2023
Author(s):	Paul CHANEY, Marc MORGAN, EBRAINS (P1)
Compiled by:	Paul CHANEY, Marc MORGAN, EBRAINS (P1)
Contributor(s):	
WP QC Review:	Anna LÜHRS, JUELICH (P20)
WP Leader / Deputy Leader Sign Off:	Steve FURBER, UMAN (P63)
T7.4 QC Review:	Annemieke MICHELS, Martin TELEFONT, EBRAINS (P1)
Description in GA:	Final release of the Collaboratory at TRL 9 level
Abstract:	This Deliverable summarises the final release of the Collaboratory at Technology Readiness Level 9. It includes brief descriptions of the latest status with regards to all related services, i.e. IAM, Lab, Wiki, Drive, Bucket and Office. In addition, an overview of related trainings is provided and supporting services are described. Also, the migration process from Collaboratory 1 to the current Collaboratory 2 is outlined. A short outlook for potential future activities is given. Links to the Collaboratory code base are provided in an Annex.
Keywords:	Collaboratory, Collaboration, Keycloak, IAM, Wiki, Xwiki, Drive, Seafile, Seadrive, Lab, Juptyer, Hub, Bucket, Data-proxy, collab, collabs, storage, team, teams, units, groups, subunits
Target Users/Readers:	computational neuroscience community, computer scientists, Consortium members, general public, HPC community, neuroimaging community, neuroinformaticians, neuroscientific community, neuroscientists, platform users, policymakers, researchers, scientific community, students

Table of Contents

Table of Tables

Table of Figures

1. Introduction

The EBRAINS Collaboratory[1] provides core services to other EBRAINS services. It also provides end users with a programming environment for reproducible science.

The core services offered in the Collaboratory are listed below; see Section 2 for a more detailed description.

- **IAM**: a single sign-on solution for EBRAINS services so that users have a single account and so that services can refer to common user IDs
- **Lab**: offers end users a programming environment for reproducible science. The Lab also offers tool developers to pre-install their tools via official EBRAINS releases so that end users do not have to go through the tedious process of installing tools and resolving compatibility issues.
- **Storage**:
  - **Drive**: offers a user-friendly UI for managing files in the cloud
  - **Bucket**: offers a simple interface to object storage optimised for larger files (e.g. datasets, videos)
- **Documentation** and more:
  - **Wiki**: offers end users a convenient way of documenting their work and to make it available to the community in web form. The Wiki is also used as a framework to develop additional services such as
    - **Identity manager**: for users to search for other users with an EBRAINS account
    - **Teams manager**: to manage permissions in Collaboratory workspaces
  - **Office**: offers the collaborative edition of documents in one of the Microsoft Office standardised formats

The EBRAINS Collaboratory is the second version of the Collaboratory developed in the Human Brain Project (HBP). It builds on existing third-party software so that more services and features can be offered to the end user at a lower cost for the Research Infrastructure (RI).

The goal of the Collaboratory is to, as the name suggests, encourage and facilitate the collaboration of scientists and the collaboration of EBRAINS tool developers. It achieves this by organising work around workspaces, known as 'collabs', which are created by users. These collabs are a concept built on top of the existing third-party software to manage permissions in a consistent way and to present an integrated view of all the Collaboratory services to the end user.

All Collaboratory services are self-hosted at Fenix[2] sites. This was an important requirement from the beginning of the HBP, as several of our Partners are forbidden from working on free commercial platforms for any of their work and all of our Partners are forbidden from using them for any confidential data.

In itself, the Collaboratory is not specialised for brain research or neuroscience. The core services it deploys could be reused in other research contexts. The Collaboratory does not provide secure storage to use for personal data (as per GDPR). Therefore, the HBP put out a Secure Sensitive Data (SSD) call and the Health Data Cloud (HDC) service is consequently being developed.

The work in SGA3 concerning the Collaboratory has consisted in:

- Fine tuning the architecture of the Collaboratory
- Implementing and improving the core services of the Collaboratory (IAM, Lab, Drive, Bucket, Wiki, Office)

[1] https://ebrains.eu/service/collaboratory
[2] https://fenix-ri.eu/

- Planning and executing the migration from Collaboratory 1 to Collaboratory 2
- Interviewing users to collect feedback on the usability of the services
- Identifying and prioritising key features to better meet user expectations
- Conceiving an official release mechanism for EBRAINS tools in the Lab environment
- Identifying unique and returning users for reporting statistics to the EBRAINS Communications team and HBP reports
- Demonstrating the Collaboratory in workshops and trainings
- Improving the UI/UX experience
- Integrating the IAM login mechanism with Fenix AAI
- Keeping the base third party software up to date for security and added features
- Improving the ease of integration of services to EBRAINS
- Documenting and creating tutorials for the Collaboratory features
- Releasing Collaboratory software at Technology Readiness Level 9 (TRL9)

The Collaboratory makes EBRAINS tools more accessible to the end users. Users do not need to install the tools for themselves. They can benefit from other users' experience by finding Jupyter notebooks demonstrating the use of tools and services.

The Collaboratory also helps to promote services. Developers can integrate their services as 'community apps' which can easily be instantiated by users in their collabs.

2. Collaboratory Services

2.1 IAM

IAM[3] is the main Authorisation and Authentication system used to manage HBP/EBRAINS user accounts and services being integrated into EBRAINS. IAM manages all logins to EBRAINS services.

The EBRAINS Access Policy[4] was fine-tuned and the procedures for validating institutions and independent users were set in place. Users affiliated to an institution/company, which has not yet been validated, can request its validation by contacting the Support. This procedure is partly automated and can typically be performed on the same day. The details of the procedure are confidential for security reasons. Users not affiliated to an institution can also request an account by providing their IDs and participating in a videoconference verification. The EBRAINS Access Review Committee has the mandate to organise and oversee the procedures.

During the writing of the Specific Grant Agreement 3 (SGA3), we considered the possibility of integrating IAM with external Identity Providers. GEANT's eduGain[5] service was considered. However, we quickly realised that a large portion of HBP Partner institutions were not covered by eduGain. Furthermore, the institutions presenting the more complex validations are typically not located in Europe and as such integration to eduGain was removed from the list of higher priorities in favour of a partial automation of the validation of institutions.

IAM stores Collaboratory permissions. The Collaboratory implements permission management within collabs at various levels of granularity using Groups and hierarchical Units. Permissions in the Collaboratory are organised as Teams. A Team determines who has access to a collab. Being added to a Team of a collab grants access not just to the Wiki, but also to the Drive, Bucket and Lab of

[3] https://iam.ebrains.eu
[4] https://ebrains.eu/terms
[5] https://edugain.org/about-edugain/edugain-and-geant/

that collab as well. Other EBRAINS services also use Teams to manage permissions for their own services.

The software Keycloak[6] (KC) was chosen to implement IAM for multiple reasons, the first being that KC is an open source tool that has been thoroughly audited, and whose development is stewarded by RedHat/IBM, a leader in the UNIX software market. KC allows the utilisation of the most common authentication protocols, including OIDC (OpenID connect) which we use. As well as this, KC allows us to self-host our client, meaning we keep everything inside the EBRAINS-controlled infrastructure and no secure data are stored with external providers. KC also allows us to identify returning users, which is vital for providing usage statistics to the EBRAINS Partners to better identify how dissemination events contribute to raising awareness to EBRAINS services. KC also allowed backwards compatibility with Collaboratory 1, an important feature to consider and discussed more in the section around the migration below (see Section 3).

The European Commission (EC) requires that we report statistics on users of the EBRAINS RI including gender and country. To fulfil this requirement, IAM pushes pseudonymised metadata to a self-hosted secure logging service. EBRAINS users are also invited to provide metadata in their profile in the PLUS project management tool (see Task T7.5). The data from these two sources are then cross-referenced and fully anonymised to provide statistics to the EC regarding usage of the various EBRAINS services.

EBRAINS tools communicate with IAM via their respective OIDC clients. This allows each tool to authenticate its users and to potentially query additional user information from IAM or to request access to core resources of each user. A major achievement in SGA3 was allowing EBRAINS tool developers to manage all of their OIDC clients, or create new ones, via an easy-to-use UI as opposed to only via the API.

2.2 Lab

The Collaboratory provides a reproducible science programming environment known as the 'Lab[7]'. This service is based on JupyterLab and JupyterHub[8]. JupyterLab provides access to Jupyter notebooks in Python and R. JupyterHub manages a separate Jupyter environment for each user.

For the end user, the Lab provides a work environment where EBRAINS tools are pre-installed via both official EBRAINS releases (every few months) and experimental builds (weekly). Users do not have to work out build options and cross-tool library compatibilities; this is all provided for them. Users can choose the more cutting edge 'experimental' version of tools or the fully tested 'releases'. Official releases will remain available online to help users continue running today's notebooks in the future. The Lab runs in the cloud on the Fenix infrastructure thereby making EBRAINS tools more accessible, e.g. to students in countries with limited computer resources.

For the EBRAINS tool developer, the Lab is an unmatched opportunity to make their tools readily available to the EBRAINS user community. The Lab provides an environment for developers to run test builds of their tools and then address merge requests to the EBRAINS Technical Coordination to push updates to the official EBRAINS releases in a scalable process, which will support a large number of tools, and which is based on a standard deployment solution (spack).

The Lab Docker image is also versioned. It includes just the base JupyterLab software and its plugins. Tool developers with specific needs can request the installation of additional extensions, e.g. a UNICORE extension was added for the TVB tool.

The integration of the Collaboratory IAM with the Fenix AAI (Task T6.6) as an identity provider allows easy access of Jupyter notebooks to the Fenix infrastructure for EBRAINS users that have requested Fenix accounts and resources.

[6] https://www.keycloak.org/
[7] https://lab.ebrains.eu/
[8] https://jupyter.org/

In collaboration with Tasks T5.12 and T5.13, automated testing has been implemented to test the integrity of the Lab Docker images, the build of EBRAINS tool releases and a selected set of user notebooks. This can typically be used for sample notebooks or for notebooks for workshops/courses.

The Lab is not intended to replace the computing power of the High-Performance Computing (HPC) infrastructure, but, for notebooks requiring more than the default 2GB RAM, users can request access to a larger Lab Docker image with 8GB RAM.

The Lab service has achieved a high level of service availability, with nearly all downtime accounted for by regular maintenance and deployments that were scheduled and announced in advance. Notebooks from Collaboratory 1 were successfully migrated and are now usable in Collaboratory 2. A full list of improvements to the Lab can be found on the Wiki at the releases page[9].

2.3 Wiki

The Collaboratory Wiki[10] is the framework that helps pull the Collaboratory services together alongside other EBRAINS services. The main way the Wiki does this is by providing the UI and the API to the workspaces called collabs.

Furthermore, the Wiki provides access to 'Community Apps', which is an easy way for other services to integrate into the Wiki.

Service providers can add their service to the list of Community Apps. Users can then easily locate these apps and add them to the navigation of their collabs. In this way, collabs can be created that will allow EBRAINS users to publish their findings alongside the services/apps they have used to create those findings.

Full version tracking is implemented in the Wiki, with users being able to view the history and to revert to a particular version of a Wiki page. Users can request the highlighting of their collabs to other EBRAINS users with the 'Highlighted collabs' feature, which will display their collabs on the home page of the Collaboratory Wiki as well as on the collab search page.

The Wiki supports the Identity Manager app that allows users to browse all groups and units available to a user as well as to search for other users. If users have the necessary rights, this is also where they can create new groups.

The Wiki also supports the Team app where collab admins manage permissions for team members within a given collab. Using either the API or the UI, people can add EBRAINS users, groups and units to a team.

More than this, we have achieved great strides in improving the Wiki for use by EBRAINS users, including large improvements to the UI, functionality and UX experience of the Wiki via numerous updates. Some of these updates include a new 'Like' feature and searching collabs with various filters like collabs a user has liked, admin/editor/viewer status, etc. Collab creation has also undergone an update and is now cleaner and easier to understand. The tour feature has been removed in favour of revamped documentation and tutorials provided in the Collaboratory collab[11]. The overall achievements are too many to recount here, but a full list can be found on the Wiki at the releases page[9].

The Collaboratory Wiki is based on the XWiki third-party software.

[9] https://wiki.ebrains.eu/bin/view/Collabs/the-collaboratory/Releases/
[10] https://wiki.ebrains.eu/
[11] https://wiki.ebrains.eu/bin/view/Collabs/the-collaboratory/

2.4 Drive/Bucket

The Drive[12] and Bucket are storage solutions offered by the Collaboratory, but both have important features that set them apart.

Drive:

- Used for mounting a file system to provide persistent storage in the Lab
  - This allows users access to all of their data and notebooks stored on their Drive
- Used for 'hot storage', for files that are changed often such as Word or Excel files
- Used for storage of smaller files with in-built editors in the Collaboratory, e.g. markdown
- Used for online editing of certain files in conjunction with the Collaboratory Office
- Full history of a file is available for viewing and/or reversion
- Based on the third-party software SeaFile[13]

Bucket:

- Used to provide object storage by allowing EBRAINS users to access Fenix Swift storage without requiring a Fenix allocation. The Bucket service is also referred to as 'data-proxy' at times.
- Used for 'cold storage', i.e. files that tend not to change as often such as videos and data-sets
- Used for storage of larger files
- Custom made software by the Collaboratory team

Drive and Bucket allow users to store large amounts of data and many types of files online without needing to set anything up themselves. They also facilitate collaboration due to how simple the storage systems allow file sharing. Both solutions allow the user to create public links for easy sharing of data with other users. Drive and Bucket also provide links to be generated that will allow anyone with the link to view the file, regardless of whether they have an account or not. The Drive also encourages easy collaboration with users outside of EBRAINS, as it is possible to generate a public upload link, which allows anyone to upload data to a particular collab without being able to view the contents of that collab.

Drive and Bucket together make up a resilient storage solution that can be integrated into most other EBRAINS services. The storage solutions also provide an API so users can fetch their data from anywhere they need. The Drive is also vital to the operation of the Lab due to the points mentioned above.

Various updates to the Drive have also brought many new features, as well as an improved UI. We also hugely improved the Lab-Drive integration, allowing many more Lab instances to run simultaneously without impacting the overall performance. We greatly reduced the loading time of a Lab session by adding loading priority for the Drive, which allows the Lab to load the Drive of the opened notebook first.

The Bucket was not included in the SGA3; however, it became obvious during SGA3 that requiring data providers to request Fenix allocations was a significant barrier for users. This led to the development of the Bucket service. Great strides have been made since the initial implementation of the Bucket in not only the functionality of the Bucket, but also to its UI. It is now a fully integrated service in the Collaboratory and is in use with multiple other EBRAINS services.

[12] https://drive.ebrains.eu/
[13] https://www.seafile.com

2.5 Office

The Collaboratory Office is integrated directly into the Drive and allows online editing of Microsoft Office files. This allows the easy collaborative editing of documents by many people simultaneously.

The Office will display, with some limitations, the user that has made the changes to a document. The limitations in question refer to the fact that, if multiple users are editing a document at once, then only one of these users will appear as an editor in the storage history, even if it was not that user that actually modified the document.

The Collaboratory Office service is based on the third-party OnlyOffice[14] software.

2.6 Collaboratory training

As part of the SGA3, the Collaboratory team was to host several training sessions (see Table 1), which has been achieved with multiple workshops and training sessions organised for various EBRAINS users and teams. Some of the trainings have also been recorded and stored online, such as a Collaboratory workshop on 25 June, 2021 that had 38 participants and which is now shared online[15].

Table 1: Collaboratory workshops and trainings

Date	Number of Participants	Event Title
5 Dec 2020	24	HBP presentation 'Collaboratory for Coordination'
13 Jul 2020	Unknown	FENS 2020, 'The Collaboratory, a gateway to EBRAINS'
24 Nov 2020	Unknown	SfN Connectome 2021, 'The Collaboratory, a gateway to EBRAINS'
13 Jan 2021	Unknown	SfN Connectome 2021, 'The Collaboratory, a gateway to EBRAINS'
15 Apr 2021	17	Migrathon
25 Jun 2021	38	Workshop
26 Jul 2021	18	Migrathon
27 Jul 2021	14	Migrathon
30 Jul 2021	13	Migrathon
23 Aug 2021	16	Migrathon
1 Sep 2021	16	Migrathon
6 Mai 2022	10	Presentation to the PHRASE project consortium
30 Mai 2022	Unknown	Brain Simulation School 2022, 'EBRAINS technology for developers'
7 Jul 2022	4	Internal training
13 Oct 2022	Unknown	HBP internal meeting, 'EBRAINS Collaboratory'
8 Nov 2022	82	Presentation at the 'Simulate with EBRAINS' workshop

2.7 Supporting services

In addition to the services mentioned in previous sub-sections, several services are deployed to support the Collaboratory:

- OpenShift/OKD
  - Used on our servers as a containerised deployment and development system
  - Used for running containers/pods for Lab users at multiple Fenix sites

[14] https://www.onlyoffice.com
[15] https://wiki.ebrains.eu/bin/view/Collabs/the-collaboratory/Workshops/Workshop 25-06-2021/

- Gitlab
  - Used as our code repository
  - Used to host our pipelines for deployments and testing
  - Also used as our organisation tool where our work is categorised into 'Tickets' and 'milestones'
- Matomo
  - For our user tracking and analytics
  - Self-hosted alternative to Google analytics
  - Data is restricted to EBRAINS service providers and stored only in anonymised form
- Docker image registry
  - Our storage and distribution system for EBRAINS docker images
  - Used for the Lab Docker images
- Uptime robot
  - A tool to monitor which services are running and their uptime over the past month
- Zabbix
  - A monitoring tool for our servers, letting us know the usage and load on our servers
- ELK logging
  - We use three services working together for syslog centralisation and dashboards:
    - Elasticsearch
    - Loghost
    - Kibana
- AWX
  - Used to run batch jobs such as backups and SSL renewals
- Minio server
  - Cache for the GitLab pipelines for building tools which are released to the Lab
- RocketChat
  - A chat service available to all users with an EBRAINS account
  - This is not officially a service integrated in the Collaboratory today, but may be further integrated in the future.

2.8 Service maturity

Throughout SGA3, the work on the Collaboratory has been focused on making the service more stable and mature. Collaboratory services have counted among the largest number of users in EBRAINS since the beginning of SGA3.

The Collaboratory 2 has been fully transparent[16] with the users on all issues that were encountered since its initial rollout. The downtime that has been experienced can be linked to maintenance or malfunctions of the infrastructure (power outage, network issues, hardware failure) where the services are deployed, or to the scheduled and announced deployment of updates to the services.

[16] https://wiki.ebrains.eu/bin/view/Collabs/the-collaboratory/Announcements/

The Collaboratory has adopted ITIL change management processes in an Agile development environment. The team works with a regular two-week sprint schedule and uses multiple development, integration and production environments. Integration environments are shared with EBRAINS tool developers for integration testing. Formal pre-release meetings have been introduced more recently to review and accept updates prior to their deployment in production. The backup procedures have been reviewed and improved to better prevent any data loss in case of malfunctions.

Updates to the third-party software used in the Collaboratory are deployed to our production environment in a timely manner. We try to avoid forks in the tools we use by submitting pull requests so that improvements are adopted by the community.

For all these reasons, having gone through the HBP Innovation team's TRL checklist, and given that the 'actual system is proven in operational environment' for nearly two years now, the Collaboratory services have been evaluated at TRL9.

3. Migration from Collaboratory 1 to Collaboratory 2

Another goal specified in the SGA3 was the transition to Collaboratory 2 and deprecation of Collaboratory 1. The Collaboratory team held several 'Migrathon' events (see Table 1) throughout 2021 to inform and assist users in the transition to Collaboratory 2 and removed all access to Collaboratory 1 in September 2021. This was a large internal milestone and the migration was rewarded with zero issues being raised by users related to Collaboratory 1 no longer being accessible thanks to our coordination effort leading up to the migration. In order to support EBRAINS services which had been using Collaboratory 1 IAM, we also introduced a 'mitreid-sub' field to the new Collaboratory.IAM to provide backwards compatibility.

As part of the support provided for the migration, the Collaboratory team introduced redirection from the Collaboratory 1 to Collaboratory 2. This allowed Collaboratory 1 users to put in redirection from their Collaboratory 1 collabs to a Collaboratory 2 URL. In this way, users can ensure that access to their data and work is not interrupted by the decommissioning of Collaboratory 1, as visitors to the old version will automatically be redirected to the updated version. Further to this, any content of Collaboratory 1 that is not explicitly redirected to a specific Collaboratory 2 URL will be automatically redirected to a generic page on the Collaboratory 2 that explains the decommissioning and redirection.

4. Looking forward

All goals specified in the SGA3 have been met. We will continue discussing with stakeholders regarding new features and improvements that can be integrated into the Collaboratory. Some of these are listed as follows:

- Analyse whether the storage solution could be unified to provide the user with a single user interface for transferring and managing files.
- Unify the search UI (collab search, wiki page search) to improve the user experience (UX) of the Wiki.
- Further integrate a chat system in the collabs.
- Tighter integration with existing EBRAINS services such as the Health Data Cloud (HDC, Task T6.15).
- Add guidance to users in their work based on the type of data they are working with and tasks performed.
- Optimise the performance of parts of the Collaboratory, e.g. by removing the dependence of linking a collab Team to a Wiki instance of a collab which would allow EBRAINS services to create

a Team without creating the Wiki and Drive instances, saving resources, speeding up the creation process and removing the dependencies which are potentially not required.

Annex A: Collaboratory Code Base

This Annex lists the repositories of the Collaboratory code base. All these repositories are public, except those for which we are in the process of assessing if making them public may be a security risk.

- Collaboratory IAM:
  - Fork:
    - https://github.com/HumanBrainProject/keycloak

      The Collaboratory fork of Keycloak
  - Extension:
    - *This repository is not (yet) publicly available. A security audit and code review are ongoing to ensure that there is no security risk in making them public.*

      The Collaboratory extension for Keycloak, adds custom themes and backwards compatibility for Collaboratory 1 IDs
- Collaboratory Lab:
  - JupyterHub:
    - https://github.com/HumanBrainProject/clb-s2i-jupyterhub/tree/s2i

      Builds a JupyterHub image for Openshift/OKD. Custom build process that add custom extensions on top of the official JupyterHub and jupyterhub-Kubespawner

      https://github.com/HumanBrainProject/clb-authenticator

      JupyterHub extension to allow refresh of IAM tokens

      https://gitlab.humanbrainproject.org/villemai/kubespawner-seadrive

      Extension to link JupyterHub with the Drive

      https://github.com/HumanBrainProject/jupyterhub-access-token-service

      Extension to link JupyterHub with IAM
  - Jupyter Lab:
    - Current Docker image for JupyterLab:

      This repository is not (yet) publicly available. A security audit and code review are ongoing to ensure that there is no security risk in making them public.

      Builds a lighter-weight Docker image in GitLab which includes JupyterLab, selected extensions for it, and some EBRAINS customisation

      https://github.com/HumanBrainProject/clb-nb-utils

      Module to make the authentication token available to the user within the notebook
    - Extensions:

      This repository is not (yet) publicly available. A security audit and code review are ongoing to ensure that there is no security risk in making them public.

      Various Lab extensions (environment variables from URL, shareable link, etc.)
    - Pipelines:

      This repository is not (yet) publicly available. A security audit and code review are ongoing to ensure that there is no security risk in making them public.

      Pipelines related to JupyterLab for build and deployment
- Collaboratory Wiki:

- Fork:
  - https://gitlab.humanbrainproject.org/HumanBrainProject/xwiki-distribution-collaboratory

    The Collaboratory fork of XWiki

- Collaboratory Bucket:
  - *This repository is not (yet) publicly available. A security audit and code review are ongoing to ensure that there is no security risk in making them public.*

    The code base for the Bucket

- Collaboratory Drive:
  - *This repository is not (yet) publicly available. A security audit and code review are ongoing to ensure that there is no security risk in making them public.*

    Extensions of Seahub (Seafile UI) for the Collaboratory, including OnlyOffice and IAM integrations
  - https://www.seafile.com/en/download/

    Official release of SeaFile used as is

- Collaboratory Office:
  - *This repository is not (yet) publicly available. A security audit and code review are ongoing to ensure that there is no security risk in making them public.*

    Mirror of the OnlyOffice releases

- Documentation:
  - https://wiki.ebrains.eu/bin/view/Collabs/the-collaboratory/
  - https://wiki.ebrains.eu/bin/view/Collabs/demo-collab