EUCAIM
CANCER IMAGE EUROPE

Project title: European Federation for Cancer Images

Project acronym: EUCAIM

Grant Agreement: 101100633

Call identifier: DIGITAL-2022-CLOUD-AI-02

D4.6 Final Federated Core services

Author(s):	Ignacio Blanquer (UPV), Alexander Harms (Health-RI), Hakim Achterberg (Erasmus MC), Enola Knezevic (DKFZ), Esther Bron (Erasmus MC, Health-RI), Diego Silveira (ITI), Konrad Lang (BBMRI-ERIC), Kurt Majcen (BBMRI-ERIC), Eman Elsayed (EGI), Laure Saint-Aubert (MEDEX).
Reviewers (WP)	WP2, WP6, Celia Martin (QUIBIM), Sorina Pop (CNRS), Diego Jimenez Fortuny (QUIBIM)
Date of delivery:	30.06.2025
Version:	V0.1
Due date:	Month 30
Type:	Demonstrator
Dissemination level:	Public

Table of contents

Abbreviations and terms

Terms	Definitions
AAI	Authentication and Authorisation Infrastructure
AARC	Authentication and Authorisation for Research and Collaboration
AI4HI	Artificial Intelligence for Health Imaging
API	Application Programming Interface
Auth	Authentication
BBMRI-ERIC	Biobanking and BioMolecular resources Research Infrastructure - European Research Infrastructure Consortium
DBMS	Database Management Service
DCAT-AP	Data Catalogue vocabulary Application Profile
EduGain	Global interfederation service that interconnects multiple identity federations
ELK	Elasticsearch, Logstash, Kibana
EMX2	Current version of Molgenis catalogue
EUCAIM	EUropean Federation for CAncer IMages
FAIR	Findable, Accessible, Interoperable, Reusable
FDP	FAIR Data Point
GDPR	General Data Protection Regulation
GUI	Graphical User Interface
IdP	Identity Provider
LS-AAI	Life Sciences Authentication and Authorisation Infrastructure
Negotiator	BBMRI-ERIC service for structured negotiation for biomedical resources
RDF	Resource Description Framework
SPE	Secure Processing Environment
SU	Support Unit
Vault	Hasihcorp identity-based secrets and encryption management system
VO	Virtual Organisation
VRE	Virtual Research Environment
WP	Work Package
XNAT	Extensible Neuroimaging Archive Toolkit

Disclaimer

The opinions stated in this report reflect the opinions of the authors and not the opinion of the European Commission.

1. Introduction

This report complements the information presented in four demonstration videos and jointly constitutes the deliverable *D4.6 Final Federated Core Services, which is an updated version of D4.5 First Federated Core Services*. EUCAIM Core services are those that support the main functionalities of the federation, and comprise: the Authentication and Authorisation, public catalogue, federated search, data access request, monitoring and helpdesk. Additional core services will be made available for data transfering, distributed processing and the storage of data.

#	Purpose	link
1	Comprehensive presentation of EUCAIM core services including the data exploration, data searching, data request and data processing on a Secure Processing Environment.	https://u.i3m.upv.es/zzh7a
2	Show data set exploration in the newest version of the catalogue, perform a request of a dataset stored in the Health-RI / Euro-BioImaging Reference Node and explore the data.	https://u.i3m.upv.es/qd7i4
3	The RDF (Resource Description Framework) endpoints and the prototype of catalogue harvesting.	https://u.i3m.upv.es/hvdgs
4	The monitoring service, showing the collection of information on the status of services and resources.	https://u.i3m.upv.es/s8kc8

This deliverable is not an extensive description of the architecture, which is kept continuously updated in an online document[1].

2. Updated Architecture

No major changes have been performed to the original architecture, although components have been improved and extended, providing new functionalities. A detailed description of the architecture is provided in https://eucaim.gitbook.io/architecture-of-eucaim.

The architecture defines three tiers of data interoperability:

- Tier 1: Compliance with the metadata model for the datasets.
- Tier 2: Compliance with the data model for searching purposes, considering the use of a mediator component that translates the queries to the specific format of the data holder.
- Tier 3: Direct (through adoption) or indirect (through a mediator component) compliance with the data model for processing purposes.

These three tier levels are related to the following federation concepts:

- Tier 1: The datasets hosted by the local node are registered in the central catalogue. Ideally, this is done through the exposure of FAIR Data Points that are used by the central catalogue to harvest the dataset's metadata, although manual registration is

[1] EUCAIM Architecture: https://eucaim.gitbook.io/architecture-of-eucaim/

also feasible for tier 1. Users can explore the metadata of the datasets registered in EUCAIM's platform.

- Tier 2: The data of the federated node[2] is searchable through its local searching service, which is queried by the federated search system through a Query Mediator component that transforms the query from EUCAIM's model to the local model and vice-versa for the results. In case the local model already complies with EUCAIM's one, the mediator component is still necessary for transforming the results and utilising the network communication middleware. The users can explore the actual number of studies fulfilling the search criteria defined by the user.
- Tier 3: The federated node has a materialisation component that makes the data available to the federated processing, according to EUCAIM's model. The user will be able to run processing actions on the actual data, if the access to them is granted.

The following URLs forward to the EUCAIM Platform services

- Dashboard (dashboard.eucaim.cancerimage.eu)
- Catalogue (catalogue.eucaim.cancerimage.eu)
- Catalogue RDF endpoint (https://catalogue.eucaim.cancerimage.eu/Eucaim/api/rdf/)
- Federated Search (explorer.eucaim.cancerimage.eu)
- Negotiator (negotiator.eucaim.cancerimage.eu)
- Helpdesk (https://help.cancerimage.eu/)
- Training platform (https://training.eucaim.cancerimage.eu/)
- Monitoring service (https://monitoring.eucaim.cancerimage.eu/)
- Repository of software artifacts (https://harbor.eucaim.cancerimage.eu/)
- Hyperontology (https://hyperontology.eucaim.cancerimage.eu/)
- Reference Node at UPV (eucaim-node.i3m.upv.es)
- Reference Node at Health-RI / Euro-BioImaging (xnat.health-ri.nl)

The summary of updates is the following:

- Dashboard now includes information on the user's profile, the datasets that a user is authorised to access to, the datasets under negotiation (user's library), improved data statistics, data holder's interaction forms and helpdesk form for unauthenticated users.
- Catalogue has been upgraded to version EMX2 of Molgenis, supporting the final terms of the dataset model and their terminologies. Support of a FAIR Data Point for exposing DCAT-AP data in RDF.
- A new version of the federated search explorer with support for the "OR" operator and 27 items included as searching criteria.
- Negotiator. Full negotiation cycle implemented in the UI, advanced dataset management and improved interface customization.
- A hyperontology service deployed to reference the terms of the hyperontology and provide searching capabilities.
- A new version of the monitoring service based on the ELK v9.0.1 has been deployed. In addition, new monitoring mechanisms have been implemented to measure the real-time resources consumption of the UPV reference node and to monitor the status of the new services deployed in it.

[2] A local node that supports federated queries and potentially also federated processing, i.e. a node that is Tier 2 or above, is called a Federated Node.

- UPV Reference node fully functional with QP-Insights® as data ingestion and dataset management tool. Integration with a secure processing environment with10 nodes, 960 cores, 7,5TB RAM and 25 GPU accelerators and an additional server with 300 TB.
- A secured registry of artifacts with three projects including public tools, tools for data preparation and tools for data processing.
- A training platform based on Moodle.
- New groups in the LS-AAI to manage authentication and authorisation of users in the Dashboard and SW registry

More details on each one of the updates are provided in the following sections.

3. Description of the Components

This section briefly describes the updated functionality and structure of each one of the components of EUCAIM described in the previous section.

3.1. Authentication and Authorisation Infrastructure (AAI)

EUCAIM Services use LS-AAI[3] as the Authentication and Authorisation Infrastructure. EUCAIM preferred this approach, rather than having a separate AAI service (e.g. a Keycloack instance) for the Central Hub services as some authorisation information is directly stored in LS-AAI.

Identifier ↓	Name	Description	ClientID/EntityID
4114	EUCAIM Harbor	The EUCAIM harbor is the repository for the certified software containers that will be used by the data holders and data users of EUCAIM to prepare and process their data.	8b6bb73f-0dc0-439a-9ec5-ad9978fd0deb
4041	EUCAIM Training service	Moodle instance for training in EUCAIM	4ffff228-5d0e-4f06-802e-46d6089140fd
4030	EUCAIM-NODE Authentication Service	The service which authenticates users in order to use any of the other services of EUCAIM-NODE platform.	5d55b0ba-aeb5-45a4-9340-18dd2263fb0a
4013	EUCAIM Explorer	EUCAIM Explorer	f0267e3e-a9de-430e-902f-51a0076aec09
3959	EUCAIM Dashboard	This is the service of the Cancer Image Europe Initiative, which provides a robust, trustworthy platform for researchers, clinicians, and innovators to access diverse cancer images, enabling the benchmarking, testing, and piloting of AI-driven technologies.	1f4a113a-1cbe-427b-8d3c-c2820ec7d08a
3951	EUCAIM Public Catalogue	Service Description as it will be displayed to end users This is the catalogue of data sets for the Cancer Image Europe Initiative, which provides a platform for researchers, clinicians, and innovators to browse, explore and request access to diverse cancer images, enabling the benchmarking, testing, and piloting of AI-driven technologies.	9a2a775c-5e1d-43e6-a59b-5060a0918a0c
3950	CHAIMELEON Authentication Service	The service which authenticates users in order to use any of the other services of CHAIMELEON platform.	b0c90d81-4354-48ec-a267-8a6f90b50ec3
3933	EUCAIM Helpdesk	A EUCAIM Helpdesk is service that facilitates support for the EUCAIM customers and users in resolution of their problems, technical incidents or questions related to the products and services of the EUCAIM project.	https://help.cancerimage.eu/auth/saml/metadata
3927	Negotiator EUCAIM - API	Negotiator instance for the EUCAIM project (API)	8c902458-3640-4418-94b4-8f4b7f338957
3895	Negotiator EUCAIM - UI	Negotiator instance for the EUCAIM project (UI)	a15a95d1-b251-4f12-b608-76ec02c68010

Figure 1: List of EUCAIM services registered in LS-AAI[4].

[3] https://lifescience-ri.eu/ls-login/
[4] EUCAIM Services management page in LS-AAI: https://services.aai.lifescience-ri.eu/spreg/auth/facilities/myServices

There is a main EUCAIM VO Group[5] (see figure 1), plus several subgroups to define specific authorisation permissions:

- Harbor: This subgroup is used to manage the push and pull permissions in the software artifacts repository of EUCAIM.
- Negotiator: This subgroup contains information about the managers of the different datasets in the negotiator.

The following services are registered in the LS-AAI (both production and development endpoints are managed by the same service identifier):

- Dashboard, deployed at dashboard.eucaim.cancerimage.eu (production) and dashboard-eucaim.grycap.i3m.upv.es (development).
- Federated Search, deployed at explorer.eucaim.cancerimage.eu (production) and explorer-eucaim.grycap.i3m.upv.es (development).
- Negotiator UI, deployed at negotiator.eucaim.cancerimage.eu. (production) and negotiator-eucaim.grycap.i3m.upv.es (development).
- Negotiator backend, deployed at negotiator.eucaim.cancerimage.eu/api (production) and negotiator-eucaim.grycap.i3m.upv.es/api (development).
- Helpdesk, deployed at help.cancerimage.eu (production).
- Reference Node at UPV, deployed at eucaim-node.i3m.upv.es (production).
- SW Artifact registry, deployed at harbor.eucaim.cancerimage.eu (production).

At this point, we have 156 users registered in the EUCAIM VO.

3.2. Dashboard

The Dashboard is an integrative application that provides an entry point for the users of EUCAIM central hub services. The dashboard has four main functionalities:

- Provide information on how different profiles get engaged with the platform.
- Collect information from the data holders through structured forms and provide this to the users.
- Host the user's library.
- Provide links to the application services.

The new main developments since last deliverable are the following:

- Integration of a helpdesk page for unauthenticated users. This feature enables users that have not yet created an account (or have problems with this) to contact EUCAIM through a form that generates a ticket in the helpdesk system (see figure 2-1).
- Integration of the forms "Data Warehouse Maturity Questionnaire" (https://dashboard.eucaim.cancerimage.eu/data-warehouse-maturity-questionnaire) and "TIER Maturity Level Questionnaire" (https://dashboard.eucaim.cancerimage.eu/tier-maturity-level-questionnaire), (see figure 2-2).
- Development of the User's library, a component to show the datasets that the logged user has been granted access to, and those which are under evaluation. This user's

[5] Enrollment URL for the EUCAIM VO Group https://signup.aai.lifescience-ri.eu/fed/registrar/?vo=lifescience&group=communities_and_projects:EUCAIM

library comprises a front-end component that queries the negotiator using the user's credentials and aggregates the information from the negotiations that have been accepted but not concluded (see figure 2-5).

- Provide a user's profile information, including the entity id in the case of Data Holders. This information is key for the CTP anonymizer tool.
- Improved information on the subjects and studies available in the platform at Tier 1 and Tier 2, aggregated by federated node (see figure 2-3 and 2-4). This information is dynamically retrieved.

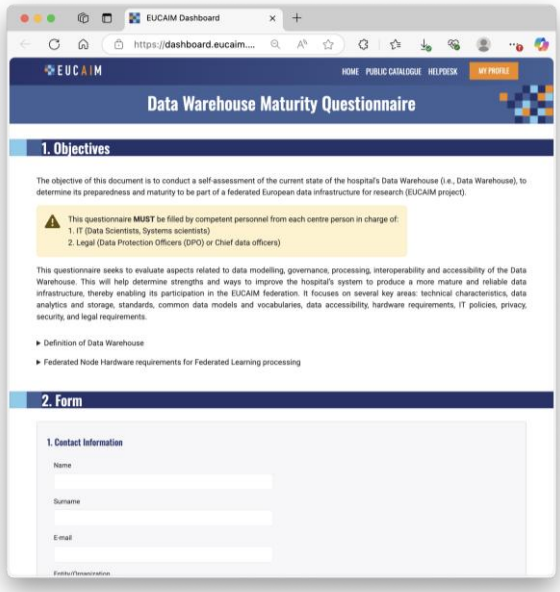

(1)

(2)

(3)

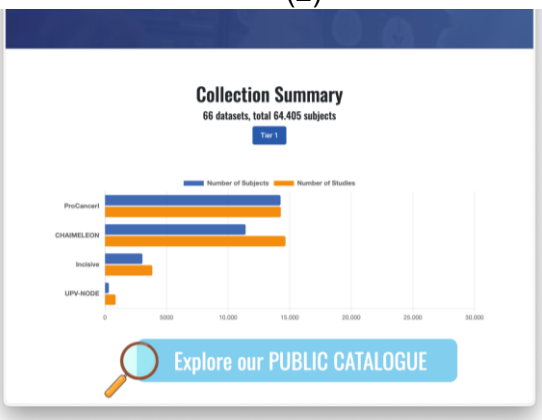

(4)

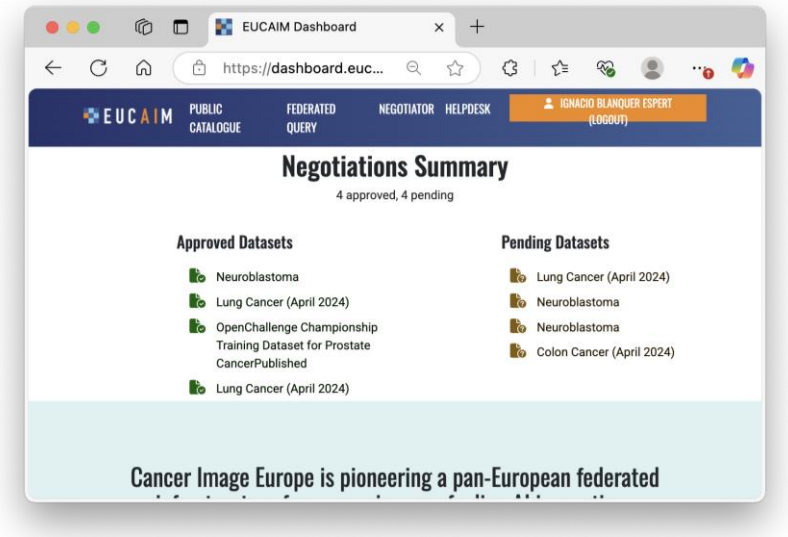

(5)

Figure 2 - New features added to the Dashboard: (1) User engagement questionnaires; (2) Helpdesk form; (3) Tier 1 and (4) Tier 2 statistics; (5) User's library.

3.3. Public Catalogue

Currently the public catalogue contains 64 datasets with a total of 45573 subjects. Several updates were made to the public catalogue. Most importantly, the catalogue was updated to the new version of Molgenis, EMX2. With this upgrade we stay aligned with the progress Molgenis has made in the development of their cataloguing software. With this major version update, the front-end also needed to be upgraded. As was done with the previous version of the catalogue, we again gratefully make use of the work that was put into the BBMRI Directory catalogue, by forking and adapting the open source code for their user interface. The interface was extended with filter options that are useful for EUCAIM users. This results in a new, refreshed look of the website (see Figure 3). Upon release of the new catalogue the internal metadata model used was also updated to match the new version of the EUCAIM metadata model. In EMX2 it is now possible to much easier set up a FAIR Data Point (FDP) by annotating the internal metadata model with the property names of the application profile used in the FDP endpoint inside the Molgenis software itself. This in contrast with the previous version, where some additional, external scripting had to be used.

EUCAIM supports two ways for DH to register a dataset in the public catalogue. The first one is a manual process supported via the Helpdesk. The process for onboarding data in the public catalogue was updated and is described in the handbook: https://docs.google.com/document/d/1Yppj9hubJ80cELR3A92k3_EOdrivSuZi68OQ1_yaHfM/edit?tab=t.0

The second way for DH to register a dataset in the public catalogue is by setting up a FAIR Data Point (FDP), from which the metadata is then harvested and added to the catalogue. The process for DH is described in *D5.8 Set-up of Local Nodes for Data Federation*. FDPs have been implemented for the reference nodes as well as several AI4HI projects, e.g:

- UPV Reference node: https://eucaim-node.i3m.upv.es/dataset-service/api/fdp/dataset
- Health-RI / Euro-BioImaging Reference node: https://fdp.healthdata.nl/
- CHAIMELEON: https://chaimeleon-eu.i3m.upv.es/dataset-service/api/fdp/datasets
- Pro-cancer-I: https://metadata.procancer-i.eu/api/fdp/
- EUCanImage: https://molgenis.eibir-edc.org/api/fdp

Currently, we are developing the FAIR Data Point harvester in collaboration with Molgenis, as this was not a standard feature of Molgenis yet. From the EUCAIM catalogue an RDF endpoint, functionally similar to an FDP, was set up. As a proof-of-principle this endpoint has been harvested into a test-setup of the EUCAIM catalogue. The code is available at https://github.com/Health-RI/molgenis-fdp-harvester. After some final alignment steps, it will be possible to harvest the CHAIMELEON FDP endpoint using the same code. Development to support harvesting from an FDP using EUCAIM-AP in the FDP reference implementation is also in progress. In D4.11 due M36, we will demonstrate that 1 or more reference node and/or project-specific FAIR data points can be harvested and their metadata can be shown in the EUCAIM public catalogue.

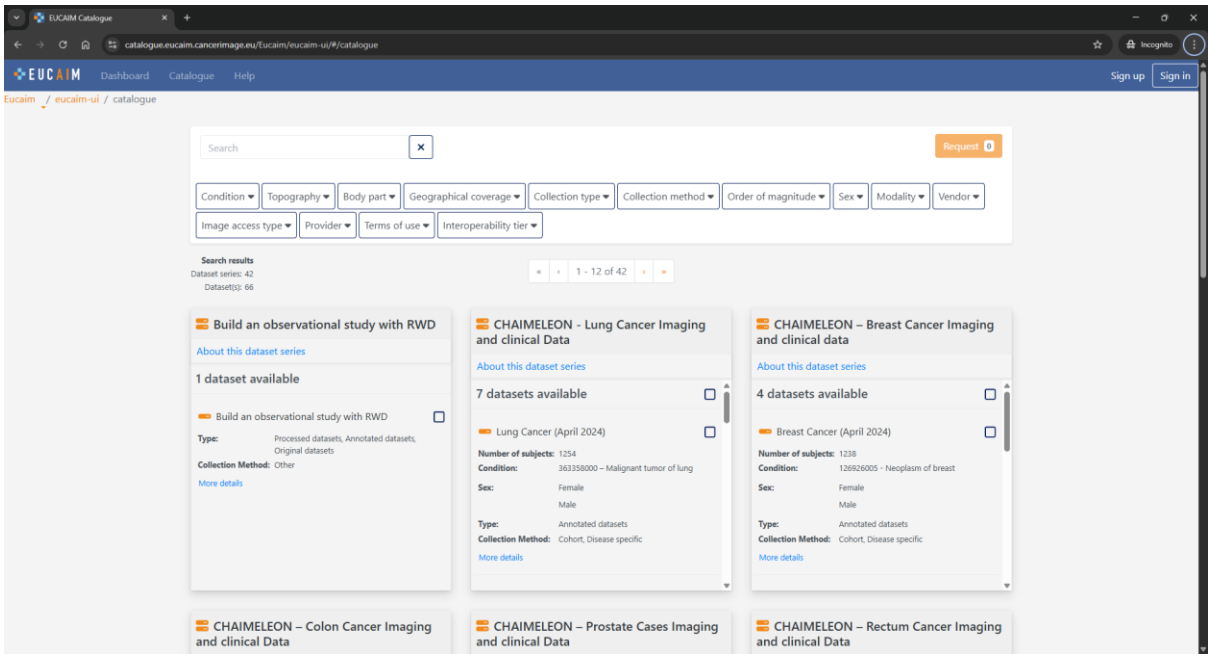

Figure 3*: The EMX2 version of the public catalogue. Unlike the previous version of the catalogue, datasets are now represented as cards that can be clicked to extend the information.*

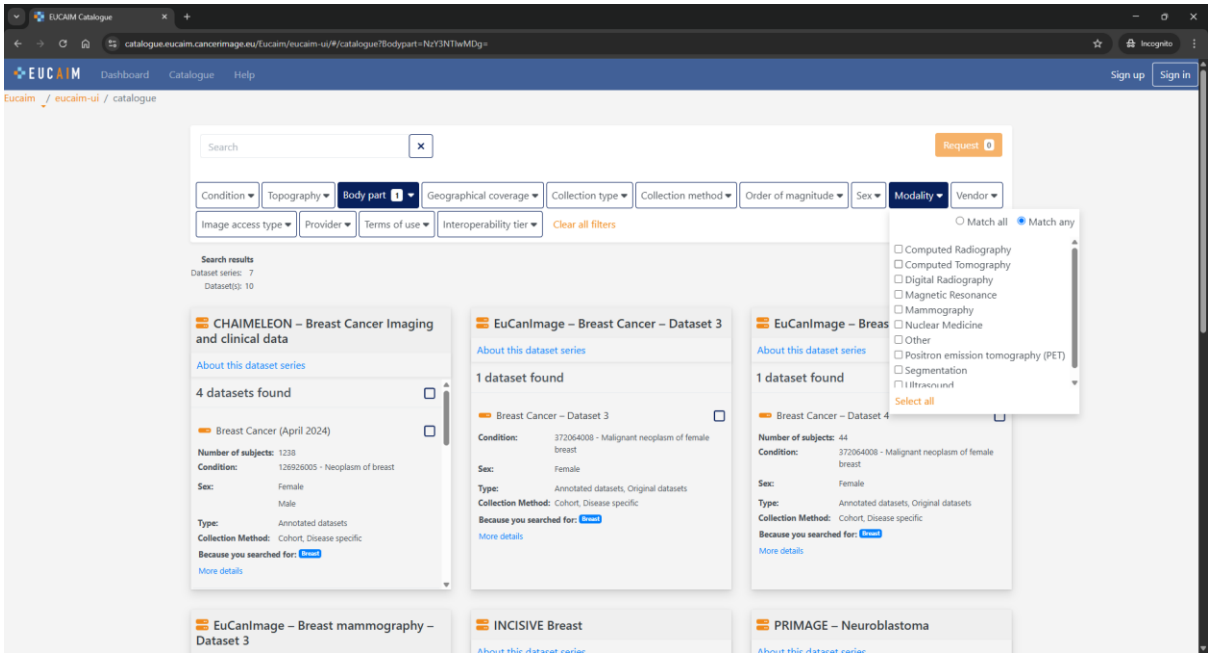

Figure 4*: The EMX2 version of the public catalogue, showing that datasets can be filtered on several variables, i.e.: Condition, Topography, Body part, Geographical coverage, Collection type, Collection method, Order of magnitude, Sex, Modality, Vendor, Image access type, Provider, Terms of use and Interoperability tier.*

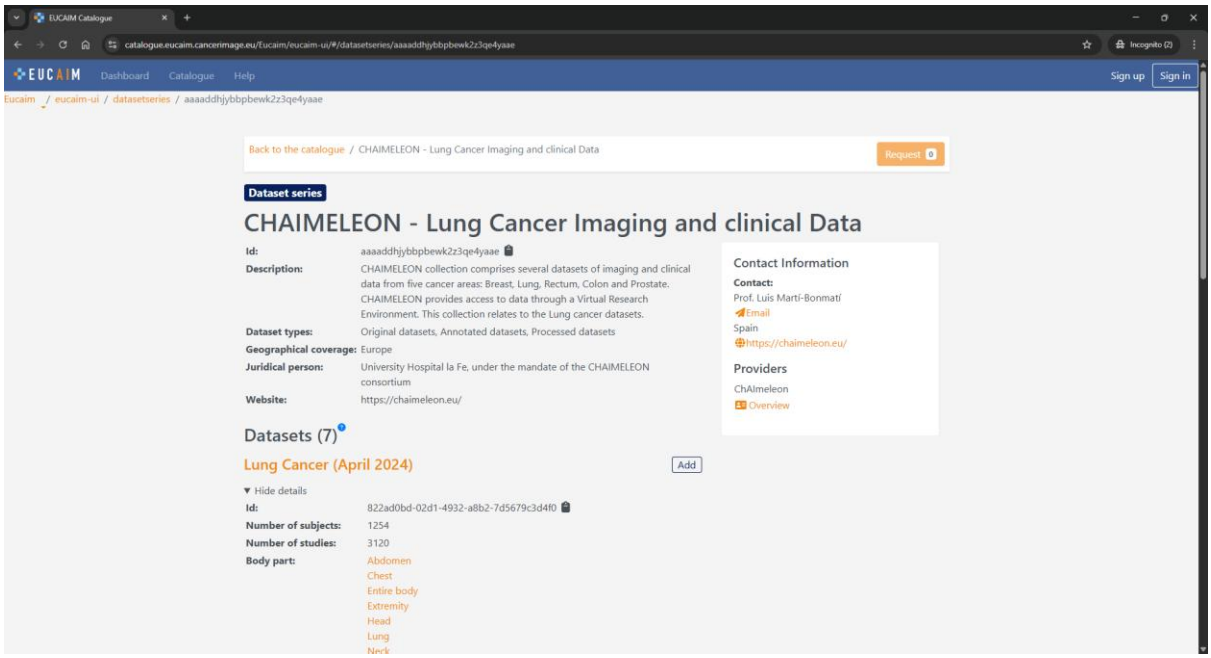

Figure 5*: Extended information in the public catalogue for one of the datasets.*

3.4. Federated Search

The federated search is a federated application based on Samply[6] technologies that consists of four main components:

- Beam is the backbone of the communication infrastructure. Beam acts as the secure message broker that connects all participating systems. It ensures that queries and responses are transmitted reliably and end-to-end encrypted between different sites without exposing sensitive data.
- Focus is the local query dispatcher at each site. It receives incoming queries from Beam and routes them to the appropriate local data sources or applications. Once the query result is retrieved, Focus sends the results back through Beam.
- Lens is the user-facing component. It provides the web interface of the Federated Search Explorer where researchers or clinicians can construct and submit queries, and explore the results. It serves as the “control panel” for interacting with the federated network and analysing the results.
- Spot works in conjunction with Beam to manage and execute queries across the federated network. It allows for the creation of new "Beam tasks" (queries) to be sent to various sites. It essentially acts as a client that sends queries to the Beam network and receives the results.

Since the initial platform release, a new version of Lens has been deployed, extending its functionality and user-friendliness. Currently 27 search variables from the EUCAIM Common

[6] https://github.com/samply

Data Model are supported for the cohort selection. Four providers are now available in the federation (see Figure 6).

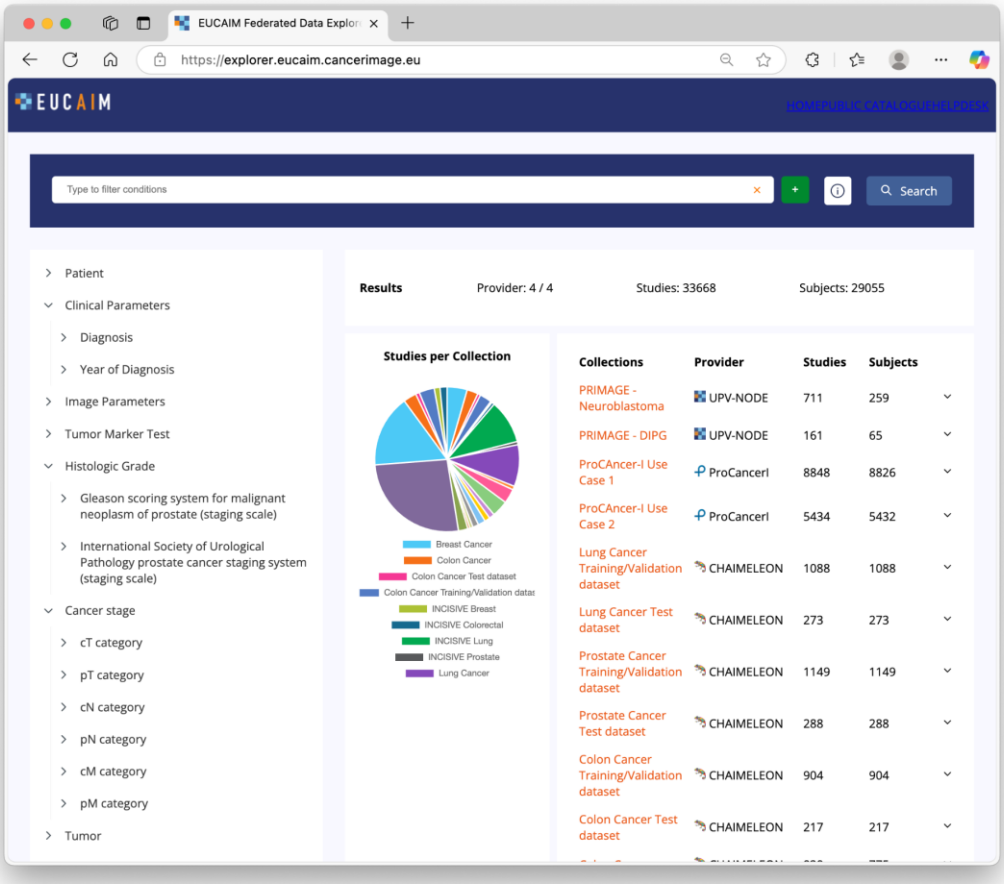

Figure 6 Federated Search Explorer, showing the new search variables (left panel) and the four providers (top bar labelled “Results” below the search bar).

3.5. Negotiator

The Negotiator is the service that manages the access requests. It is connected to the catalogue and receives the access requests to specific datasets initiated in the catalogue. The Negotiator, developed by BBMRI-ERIC[7], considers three types of actors: the requester, the dataset manager and the administrator. The current release (v3.16.4) additionally includes:

- Security updates to address CVEs detected in the cybersecurity analysis.
- Email template editing, facilitating the configuration of customized notification e-mails (see figure 7).
- Administrators can see user logins (last login date).
- Webhooks to allow for easy integration with other services.
- Information requirements can be used to use additional forms when the state of negotiation changes.
- Additional bug fixes.
- Higher level of customisation of the user interface.
- Full support of the negotiation lifecycle.

[7] https://github.com/BBMRI-ERIC/negotiator

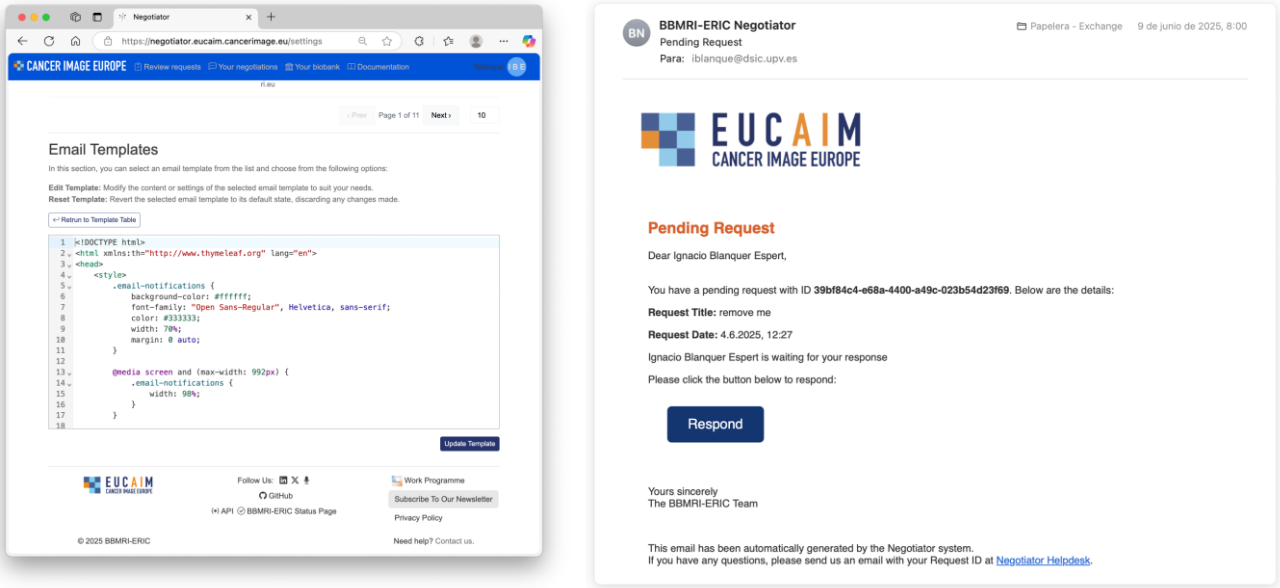

Figure 7: Mail template and mail received.

3.6. Monitoring

The monitoring service is composed of a set of tools from the ELK technology stack (Elasticsearch, Logstash, Kibana), including Heartbeat and Elastic Agent. These tools enable the collection, storage, analysis and visualization of large volumes of data in real time.

Since the previous deliverable, progress has been made in several areas:

- Version 9.0.1 of the Elasticsearch, Elastic Agent, Kibana, and Heartbeat services has been deployed.
- New monitoring mechanisms have been implemented to measure real-time resources consumption of the UPV reference node.
- A new index has been defined in Elasticsearch to store the data collected by the Elastic Agent service (see Figure 8).
- New visualization dashboards have been added in Kibana to represent the data gathered by the Elastic Agent service (see Figure 9 and Figure 10).
- Additional monitors have been added to measure the availability of the new services deployed in the UPV Reference Node (see Figure 11), such as: Harbor registry, brokers, proxies, etc.

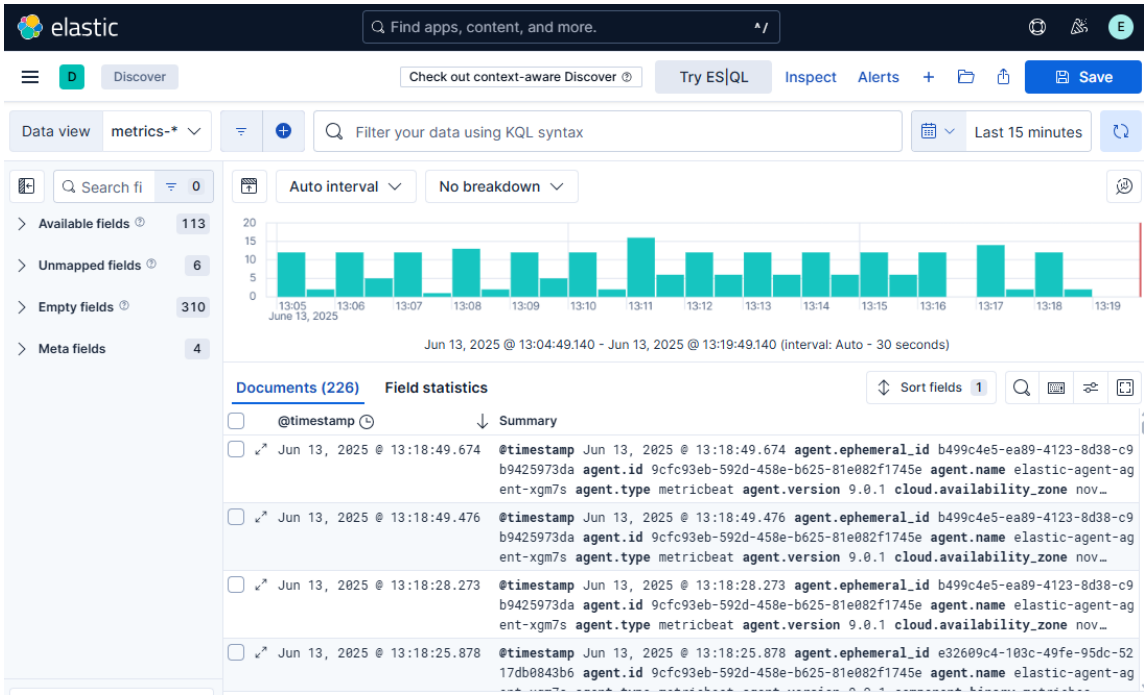

Figure 8*: Elasticsearch index storing the metrics obtained by Elastic Agent.*

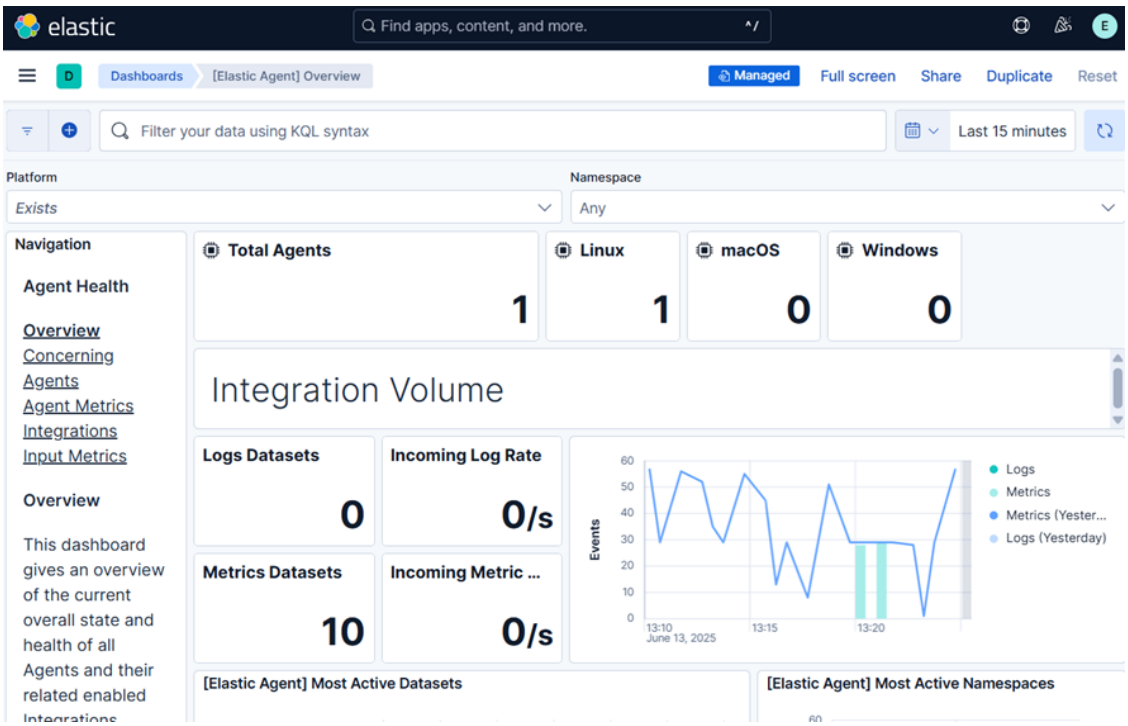

Figure 9*: Dashboard representation of the Elastic Agents Overview*

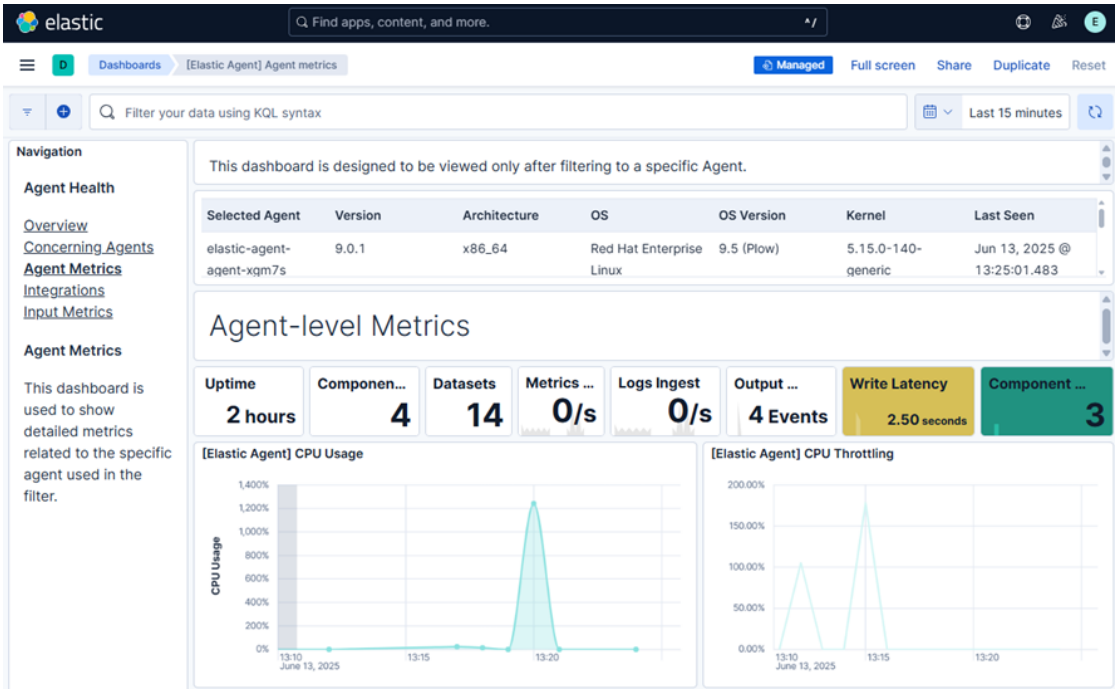

Figure 10*: Dashboard representation of the Elastic Agent Metrics*

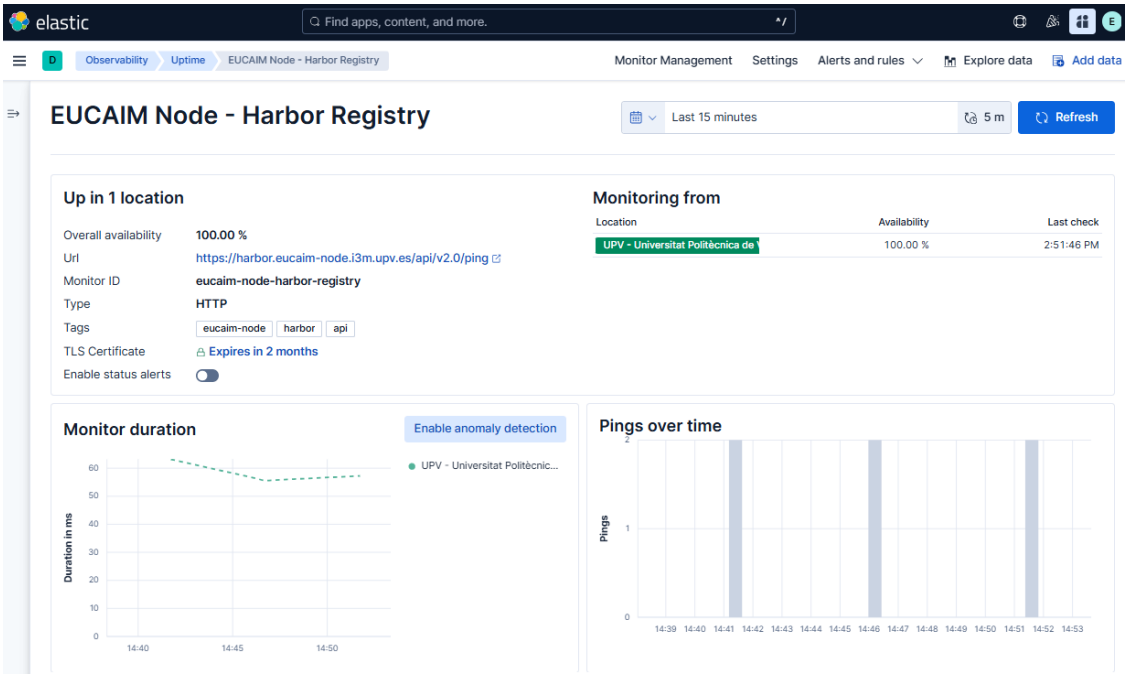

Figure 11*: Representation of the availability of Harbor registry in the UPV Reference Node*

3.7. Helpdesk

The helpdesk instance is hosted by the Karlsruhe Institute of Technology (KIT) and is configured to accept LS-AAI accounts.

It works as a ticketing system, whose goal is to record and manage submitted tickets in order to assign and reply to senders in a structured and time manner. It aims to provide assistance to any stakeholder (including consortium partners) of the project, by addressing their requests

and questions on any topic: data holder onboarding, local node setup, dataset creation, use of tools, access issues to a core service, etc.

In order to best address all types of tickets, support units have been created :

- **First line support:** default support unit when creating a ticket. Addresses all basic issues and general inquiries, and redirects the ticket to other units when relevant. If a user is not sure whom to assign a ticket to, the ticket should be assigned to the First Line Support Unit. A weekly rotation is in place for the first-line team to respond to incoming tickets.
- **Catalogue support**: addresses requests to access datasets and queries about specific datasets as well as requests to populate the catalogue based on a filled metadata excel file by the data holder.
- **Dashboard support**: addresses issues to access the webpage, and queries about how to become a data provider/tool provider/a user and, about the tools available on the platform.
- **Enrolment support**: addresses membership requests to the EUCAIM group
- **LS AAI support**: addresses issues with account creation and account login
- **Federated search support**: addresses issues related to the use of the federated search engine
- **Negotiator support**: addresses access requests and issues with the “My library” section. The helpdesk system is integrated with the **BBMRI-ERIC RT system** to streamline the management of tickets assigned to the Negotiator support unit.
- **Security support**: addresses reports on security breach and other security-related issues or suggestions.
- **Reference node support**: provides support to users of the reference nodes at UPV and Health-RI/Euro-BioImaging regarding a.o. data upload and data access.
- **Technical support**: provides support to data providers for establishing and integrating a data node, and addresses issues about dataset preparation.

Each of these support units is composed of at least one specialist, and often one or more additional technical experts. The scope of the units is clearly defined in a “Knowledge Base” environment, accessible to support unit members in the Helpdesk UI.

The EUCAIM Helpdesk is accessible:

- From the EUCAIM dashboard once authenticated with a user profile.
- From the EUCAIM helpdesk url https://help.cancerimage.eu/ (demo accessible here : https://drive.google.com/file/d/1hx8tiYI1zbXp1Vr0HyVcFI0J4OoQVhiV/view?usp=drive_link)
- By sending an email to eucaim-helpdesk@scc.kit.edu.

Dashboard and email requests will automatically land on the first-line support unit.

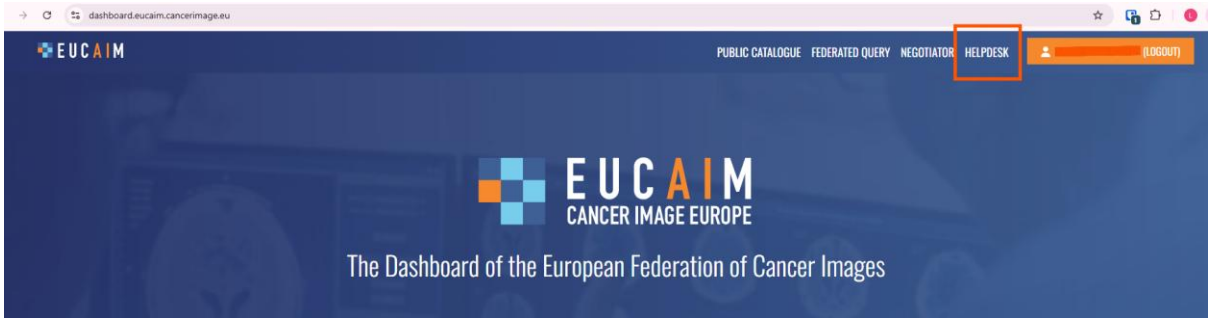

Figure 12: Helpdesk link from the Dashboard.

Deliverable D2.6 Helpdesk activity report due in December 2025 will provide insights on the feedback obtained through the helpdesk.

3.8. Software Artifacts Registry

The Platform requires managing the binaries of the applications that are validated for their execution in the preparation and postprocessing of data. The details of the application development will be described in D6.3, and this section describes the interaction with the software artifacts registry deployed.

The software artifacts registry relies on Harbor[8], and it is accessible at the URL https://harbor.eucaim.cancerimage.eu. The registry manages its authentication through the LS-AAI service, and has defined three projects and two authorisation groups. The projects are:

- **ingestion-tools**, which contains the images of the software components used in the preprocessing of the data. They could be downloaded to the data holder's premises and comprise both OCI-compliant containers and binaries encoded in an OCI envelop and accessible through ORAS[9].
- **processing-tools**, which contains the images of the tools used within the platform for the processing of the images. All the images are OCI-compliant containers.
- **library**, which contains the images of the services of the Software components in the central hub. All the images are OCI-compliant containers.

And the authorisation groups are:

- urn:geant:lifescience-ri.eu:group:lifescience:communities_and_projects:EUCAIM:harbor:**pull**#aai.lifescience-ri.eu, associated with the role "Limited Guest", which enables **pulling** artifacts from the registry.
- urn:geant:lifescience-ri.eu:group:lifescience:communities_and_projects:EUCAIM:harbor:**push**#aai.lifescience-ri.eu, associated with the role "Developer", which enables **both pulling and pushing** images in the registry.

[8] https://goharbor.io/
[9] https://oras.land/

The access to the platform requires that the users are enrolled in the harbor group[10] in LS-AAI, and belong to the pull or push subgroups[11]. This enrollment is performed the first time that a user tries to log into the harbor registry. Once accepted, the user can access the processing-tools or the ingestion-tools repositories with the appropriate permissions. The library project is readable for even anonymous users if they are provided with the digest codes of the containers[12], but only the platform administrators can push objects in the project and only authenticated users can list the objects and get the digest codes.

Figure 13 shows the list of projects and the SW artifacts in the ingestions_tools project.

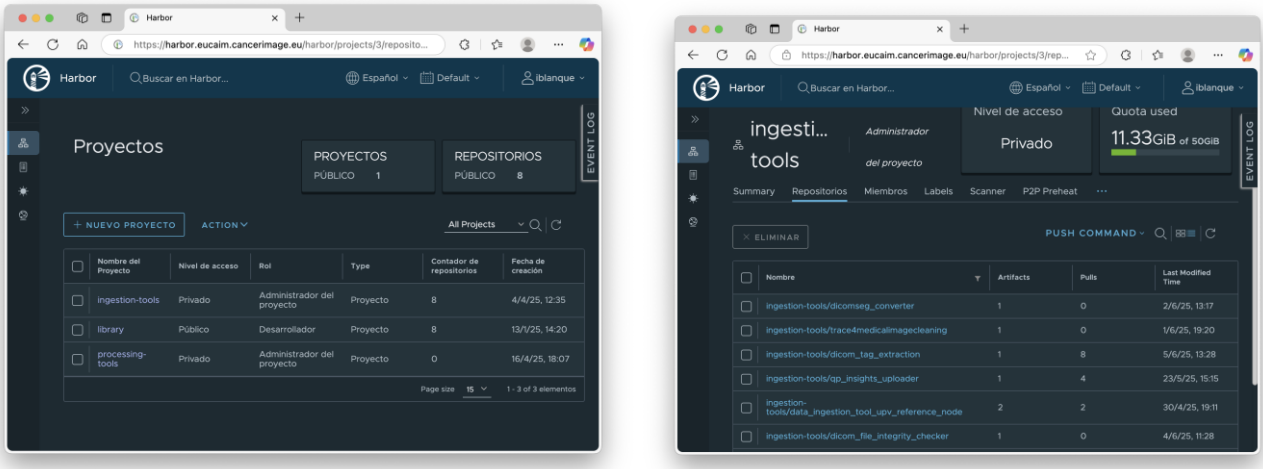

Figure 13: Software artifacts registry (harbor). List of projects (left) and binaries in the ingestion_tools project (right).

[10] https://perun.aai.lifescience-ri.eu/organizations/3345/groups/43153

[11] https://perun.aai.lifescience-ri.eu/organizations/3345/groups/43153/subgroups, note, this URL is provided as a reference, as it requires elevated permissions to access.

[12] For example, any user can download the negotiator UI container using the command: docker pull harbor.eucaim.cancerimage.eu/library/negotiator/negotiator@sha256:1212f040a2315615e94c39e95a1e34cfe4447e55154687a997eeea85902a5413 . Only authenticated users can get the digest code.

4. Outline of the Video

The demonstration videos show the process of browsing the catalogue, searching for data, selecting a dataset, requesting access, negotiating the access and accessing the data. This video complements and updates the previous video (https://dashboard.eucaim.cancerimage.eu/D4-5_7-video.mp4) developed in month 18.

The first video[13] focuses on the main services of the central hub and covers the following parts:

- Dataset exploration, which deals with the process of browsing the collections and filtering the data.
- Access Request, triggered from the catalogue and using the negotiation component.
- Access Negotiation, through the negotiator component, which may involve multiple steps.
- Access granted on one dataset stored in the UPV reference node .
- Browse the user's library to see the datasets that a user has access granted.
- Access one of the datasets in the UPV reference node.
- Explore and process the data.

The second video[14] covers:

- Data set exploration in the newest version of the catalogue
- Access Request, triggered from the catalogue and using the negotiation component.
- Access Negotiation, through the negotiator component, which may involve multiple steps.
- Access granted on one dataset stored in the Health-RI / Euro-BioImaging Reference Node.
- Explore the data.

In the third video[15] the following is shown:

- The RDF endpoint of the EUCAIM catalogue, and its configuration with semantic annotation through the configuration of Molgenis EMX2.
- The harvesting of this endpoint with the FAIR Data Point (FDP) harvester, which adds the metadata from the EUCAIM catalogue to an identical catalogue in a test environment.

The last video[16] covers the monitoring of services:

- Firstly, it shows the log files that are stored in the Elasticsearch..
- Next, the metrics collected by the different Elastic Agents on resources consumption in real time are shown in a Dashboard, including CPU consumption, throttling, RAM memory consumption, disk write metrics.
- Finally, it shows the monitoring of different EUCAIM services in real time, and what type of error has occurred in each of the checks.

[13] https://u.i3m.upv.es/zzh7a
[14] https://u.i3m.upv.es/qd7i4
[15] https://u.i3m.upv.es/hvdgs
[16] https://u.i3m.upv.es/s8kc8

5. Conclusions

This report complements the information presented in the three demonstration videos listed in the previous section and jointly constitutes the deliverable *D4.6 Final Federated Core services*. EUCAIM Core services are those that support the main functionalities of the federation, and comprise: the Authentication and Authorisation, the public catalogue, the federated search service, the data access request service, the service monitoring and the helpdesk. Additional services have been added to the platform, including the training platform, the hyperontology and the SW registry.

The deliverable shows the advances with respect to the previous version of the core services, outlining improvements:

- In the data discovery through an extended set of variables and providers.
- In the metadata exploration by the support of FAIR Data Points and the adoption of the Health DCAT-AP schema.
- In the access request through the complete management of the lifecycle in the negotiator tool
- In the Dashboard gathering information about the user's permissions and updated information.
- In the helpdesk, hyperontology service and the monitoring, providing production-ready functionality.

The information of this deliverable is complemented with additional information in the following technical documents:

- End User Guide: https://eucaim.gitbook.io/enduserguide/, which describes both the usage and building of the platform.
- Architecture guide: https://eucaim.gitbook.io/architecture-of-eucaim/
- Handbook for Data Holders: https://u.i3m.upv.es/2mtyb.