

# D6.5 HL7 FHIR Contributions

**Lead Author:** Luc Chatty

**With contributions from:** Giorgio Cangioli, Rhonda Facile

**Reviewer:** Olga Méndez, Berta Miró, Marta Wilinska

<b>Deliverable nature</b>	Other
<b>Dissemination level</b>	Public (PU)
<b>Delivery date</b>	30-04-2026
<b>Version</b>	1.0
<b>Total number of pages</b>	48
<b>Keywords</b>	HL7, FHIR, Standardization, Implementation Guide, Interoperability

## EXECUTIVE SUMMARY

The FLUTE project contributes to advancing interoperable, standards-based digital health solutions aligned with the European Health Data Space (EHDS). Within this context, HL7 Europe has played a central role in ensuring that project outputs are aligned with international interoperability standards, particularly HL7 FHIR.

HL7 Europe delivered both standardisation and implementation outcomes. This includes integrating FLUTE requirements into HL7 processes, contributing to cancer data modelling through the PHOENIX initiative, and enabling interoperability between clinical, imaging (DICOM), and research data. In addition, the project demonstrated how HL7 FHIR can support federated research workflows and reproducible AI pipelines.

Key assets developed within FLUTE include FHIR Implementation Guides (R4 and R5), reusable modelling artefacts for research and oncology, and a sandbox environment enabling end-to-end validation of interoperability workflows. HL7 Europe also supported partners through training, workshops, and technical guidance, strengthening their capacity to adopt standards-based approaches.

These contributions demonstrate that existing interoperability standards can effectively support complex use cases such as federated learning and cross-border research. By embedding FLUTE outputs within recognised European and international frameworks, the project ensures their reusability, scalability, and sustainability beyond its lifecycle, while contributing to the broader development of the EHDS.

## DOCUMENT INFORMATION

Grant agreement No.	101095382	Acronym	FLUTE
Full title	<b>Federate Learning and mUlti-party computation Techniques for prostatE cancer</b>		
Call	HORIZON-HLTH-2022-IND-13-02		
Project URL	<a href="https://cordis.europa.eu/project/id/101095382">https://cordis.europa.eu/project/id/101095382</a>		
EU project officer	Nihal Yildirim		

Deliverable	Number	D6.5	Title	HL7 contributions
Work package	Number	WP6	Title	Ethics, regulatory acceptability, GDPR guidelines and standards
Task	Number	T6.4	Title	HL7 FHIR Contribution

Date of delivery	Contractual	M03	Actual	M05
Status	version 0.5 <input type="checkbox"/> Final version			
Nature	<input type="checkbox"/> R <input type="checkbox"/> DEM <input type="checkbox"/> DMP <input type="checkbox"/> DEC <input type="checkbox"/> ETHICS <input checked="" type="checkbox"/> OTHER			
Dissemination level	<input checked="" type="checkbox"/> Public <input type="checkbox"/> Sensitive			

Authors (partners)	HL7 Europe
Responsible author	Luc Chatty
	Luc.chatty@hl7europe.org

Summary (for dissemination)	<p>HL7 Europe's contributions to FLUTE combine standardisation and implementation, including integrating project requirements into HL7 processes, advancing cancer data modelling through PHOENIX, enabling interoperability across clinical, imaging and research domains, and exploring HL7 FHIR for AI and federated research. Through training, implementation guides, workshops, and a sandbox environment, these efforts ensure that FLUTE outputs are reusable, scalable, and aligned with European and global interoperability frameworks beyond the project lifecycle.</p>
Keywords	HL7, HL7 FHIR, Standardization, Implementation Guide, Interoperability

<b>VERSION LONG</b>			
<b>Issue Date</b>	<b>Rev. No.</b>	<b>Author</b>	<b>Change</b>
<b>15-11-2025</b>	0.0	Luc Chatty	Initial
<b>19-02-2026</b>	0.1	Rhonda Facile	Added some narrative parts
<b>02-03-2026</b>	0.2	Giorgio Cangioli	Added EHDS work link and context
<b>23-03-2026</b>	0.3	Luc Chatty	HL7 contributions to partner
<b>14-04-2026</b>	0.4	Rhonda Facile	Review and harmonization
<b>16-04-2026</b>	1.0	Luc Chatty	Validation for project internal review
<b>22-04-2026</b>	1.1	Berta Miró	Reviewed and added comments
<b>23-04-2026</b>	1.2	Olga Méndez	Reviewed and added comments
<b>28-04-2026</b>	1.3	Luc Chatty	Review validation and finalizing for publication

## TABLE OF CONTENTS

### Contents

1	HL7 contribution to European and global standardisation.....	10
1.1	HL7 Europe's contribution to FLUTE and the EHDS .....	10
1.2	Cancer data standardisation – PHOENIX Initiative .....	12
1.3	Defining shareable research artefacts.....	17
1.4	Link with DICOM imaging .....	19
1.5	Standardisation of AI integration .....	20
2	HL7 contribution to partners .....	22
2.1	Training .....	22
2.2	Specification designing and sharing.....	22
2.3	Workshop with partners .....	23
2.4	Standard usage illustration .....	24
2.5	HL7 Europe Facilities - Sandbox .....	33
2.6	Future Recommendations.....	47
3	Conclusions.....	48

## LIST OF FIGURES

- Figure 1: European FHIR Implementation Guides
- Figure 2. European Cancer Common Conceptual Model
- Figure 3. Extract of the European cancer patient Journey
- Figure 4: Questionnaire response generation
- Figure 5: Pseudonymisation workflow
- Figure 6: Checking eligibility workflow
- Figure 7: Evaluation of the variable for one patient workflow
- Figure 8: Making data available in multiple formats
- Figure 9: Workflow - Creation of a cohort and a datamart for a study
- Figure 10: Workflow - Making data available for AI evaluation in care
- Figure 11: FLUTE experimental sandbox architecture
- Figure 12: CQL Engine diagram
- Figure 13: Questionnaire engine diagram
- Figure 14. Example of visualization of the FLUTE questionnaire on the sandbox
- Figure 15. Example of the CDS Hooks card component
- Figure 16. Integration of the CDS hooks card component in an App
- Figure 17: Architecture of the mapping engine
- Figure 18: Cohort engine workflow
- Figure 19: Datamart engine workflow
- Figure 20: CDS Hooks engine workflow

## ABBREVIATIONS AND ACRONYMS

AI: Artificial Intelligence  
CDISC: Clinical Data Interchange Standards Consortium  
CODEX: HL7 FHIR Accelerator for oncology and clinical specialty data standardisation  
CDS: Clinical Decision Support  
CQL: Clinical Quality Language  
DICOM: Digital Imaging and Communications in Medicine  
EU: European Union  
EUCCDM: Cancer Common Data Model  
EHDS: European Health Data Space  
FHIR: Fast Healthcare Interoperability Resources  
GDPR: General Data Protection Regulation  
HL7: Health Level Seven  
OMOP: Observational Medical Outcomes Partnership  
OMOP CDM: OMOP Common Data Model  
SDC: Structured Data Capture

## Introduction

The FLUTE project advances interoperable digital health solutions in Europe by delivering standards-based outputs that support scalable, cross-border data exchange and alignment with the EHDS. In doing so, it contributes to strengthening Europe's digital health ecosystem through the development of practical, reusable technical assets grounded in established interoperability frameworks.

Within this context, HL7 Europe has implemented a structured approach to contribute to the evolution of HL7 FHIR standards, ensuring that the interoperability needs and gaps identified within FLUTE are appropriately captured and communicated to the relevant standardisation processes. In line with the task objectives, this has included: (a) the collection and analysis of requirements and gaps identified across FLUTE work packages; (b) the establishment of liaisons with relevant HL7 Working Groups; and (c) the sharing, formalisation, and discussion of consolidated feedback within HL7 standardisation fora. Through these activities, HL7 Europe has supported the positioning of FLUTE requirements within ongoing HL7 FHIR developments, with the aim of informing future versions of the standard as well as related implementation guides and profiles.

Beyond its contribution to standardisation processes, HL7 Europe has supported the implementation and uptake of interoperability standards within the project. This includes the development of FLUTE-specific HL7 FHIR Implementation Guides and associated artefacts, as well as targeted support to project partners in applying HL7 FHIR to concrete use cases. These efforts have enabled the translation of project requirements into actionable, standards-aligned specifications, thereby supporting consistency, interoperability readiness, and early validation of project outcomes.

By combining contributions to international standardisation with implementation-oriented activities, HL7 Europe has strengthened the exploitation potential of FLUTE results and supported their sustainability beyond the project lifecycle. Embedding project outputs within recognised standards frameworks facilitates their reuse across initiatives, promotes alignment with European and international interoperability efforts, and enhances their potential for adoption in future deployments within the evolving European digital health landscape.

## 1 HL7 contribution to European and global standardisation

Within the FLUTE project, HL7 Europe has contributed to strengthening alignment between project outputs and the broader European and international standardization ecosystem. Building on HL7 standards primarily HL7 FHIR and its established governance and community processes, HL7 Europe supported FLUTE in positioning its technical work as “FAIR” (Findable, Accessible, Interoperable and Reusable) making the data, interoperable, shareable, and compatible with ongoing initiatives beyond the project scope while aligning with the evolving European Health Data Space (EHDS).

These contributions were carried out through active participation in HL7 European and International Working Group Meetings (WGMs), targeted collaborations with key standardisation communities (e.g., imaging and research), and structured engagement mechanisms such as recurring coordination meetings. This approach ensured that FLUTE project requirements and lessons learned were continuously brought into standardisation discussions, while emerging standards and best practices were channelled back into the project.

HL7 Europe’s work focused on four complementary pillars: (i) cancer data standardisation and alignment with the EU Common Cancer Data initiatives, (ii) definition of shareable research artefacts supporting federated and privacy-preserving workflows, (iii) link between HL7 FHIR and imaging standards (DICOM), and (iv) contributions to the evolving standardisation landscape for AI integration in healthcare. Together, these actions support the long-term sustainability and cross-project adoption of FLUTE outcomes, reinforcing their relevance for both European deployment and global interoperability efforts.

### 1.1 HL7 Europe’s contribution to FLUTE and the EHDS

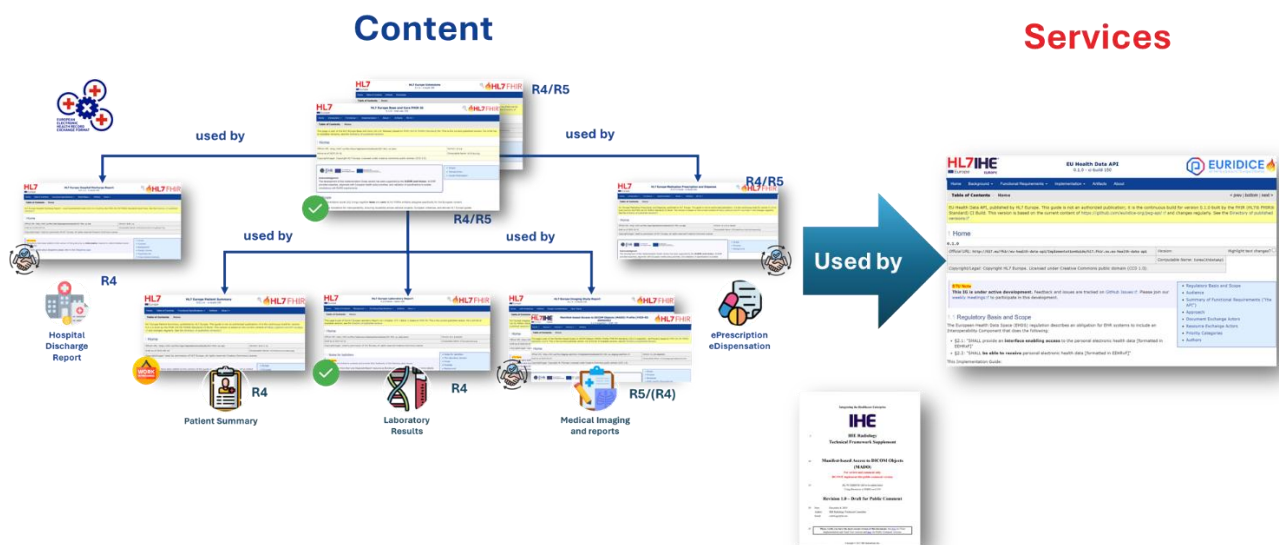
The EHDS is a key European Union regulation aimed at enabling secure, interoperable, and trusted access to electronic health data across all Member States. It places strong emphasis on patient rights and empowerment, and on the secondary use of health data for research, innovation, policy-making, and public health.

To make these rights effective, EHDS promotes the availability of standardized health data across Europe through the definition of a European Electronic Health Record Exchange Format (the “**European EHR Exchange Format**”). The EU-wide availability of this structured and semantically interoperable information will create the foundation for cross-border health data exchange, reusable digital health services, and a scalable ecosystem of interoperable solutions supporting both primary and secondary use of health data.

Although the “format” will be formally defined by the European Commission by March 2027 through a set of Implementing Acts, several preparatory initiatives already provide a reference framework that adopters can rely on:

- A reference logical model proposed by the XT-EHR common model (<https://www.xt-ehr.eu/fhir/models/>).

- A set of content-oriented European FHIR Implementation Guides published by HL7 Europe (see figure below), including a set of Core FHIR Profiles acting as reusable building blocks (<https://hl7.eu/fhir/>).
- A set of HL7 FHIR-based APIs, currently under development, to facilitate standardized access to EHDS core data (<https://build.fhir.org/ig/euridice-org/jwg-api/en/>).



**Figure 1: European FHIR Implementation Guides**

For FLUTE, the progressive availability of EU-standardised health data is a key enabler.

Harmonised clinical datasets and structured imaging-related information allow data generated in different Member States and systems to be consistently interpreted, exchanged, and reused. This is particularly relevant for clinical and imaging domains, where semantic consistency, structured metadata, and interoperable access mechanisms are essential to support cross-border collaboration, advanced analytics, and innovative clinical use cases.

Through the active involvement of HL7 Europe in European standardisation activities related to the European EHR Exchange Format and associated FHIR specifications, a continuous alignment and feedback loop has been established between FLUTE and the broader EU standardisation process. This ensures that FLUTE remains aligned with evolving European specifications, while at the same time contributing practical implementation experience and domain-specific requirements — particularly in the clinical and imaging fields — back into the European interoperability framework.

For more information on interoperability between HL7 FHIR and imaging standards (DICOM) see section below

## 1.2 Cancer data standardisation – PHOENIX Initiative

At the start of FLUTE, cancer data standardisation was not yet supported by a dedicated international HL7 work group. Existing momentum in this domain was primarily driven by the US-based HL7 Codex Accelerator, which coordinates cancer interoperability efforts across a specific ecosystem of stakeholders.

To address this gap from a European perspective and to ensure that FLUTE requirements could be anchored in a sustainable standardisation pathway, HL7 Europe initiated a dedicated cancer data standardisation task. This initiative aimed at fostering alignment across European stakeholders, supporting interoperability use cases relevant to FLUTE, and providing a structured channel to connect project outputs with HL7 FHIR modelling activities.

Through this task, HL7 Europe contributed to establishing a shared baseline for cancer-related information modelling and facilitated the integration of FLUTE-driven insights into a broader European and international interoperability context.

### 1.2.1 PHOENIX Project Start-up

To ensure that cancer data standardisation activities initiated within FLUTE continue beyond the project lifetime, HL7 Europe launched the PHOENIX Initiative as a structured, long-term European collaboration framework (<https://www.hl7europe.org/standards/#phoenix>). PHOENIX provides continuity for HL7 FHIR-based oncology modelling, governance coordination, and stakeholder engagement, enabling European partners to collectively shape interoperable cancer specifications aligned with HL7 FHIR.

The initiative focuses on testing and refining common oncology-related HL7 FHIR resources, harmonising cancer-specific profiles, and validating shared artefacts through collaborative technical work. It also promotes alignment with global accelerators such as the HL7 CodeX Accelerator (<https://codex.hl7.org/>), helping to avoid divergent modelling approaches between regions.

PHOENIX brings together key European reference institutions—including IKNL, INCa, and Charité, as well as additional partners involved in FLUTE and other related initiatives. Through regular coordination meetings and joint modelling sessions, participants align priorities, exchange implementation feedback, and converge on reusable artefacts supporting cancer data exchange, secondary use, and federated research.

By establishing PHOENIX, HL7 Europe has ensured that FLUTE's oncology interoperability outputs transition from project-specific deliverables to a durable European standardisation

trajectory, facilitating broader adoption across countries, institutions, and future EU programmes.

### 1.2.2 Common European Cancer Model

The HL7 Europe Cancer Common Data Model (ECCDM) initiative aims to define a European common data model for cancer: a minimal, shared, and extensible set of concepts and relationships capable of representing cancer-related information across different contexts and use cases.

The scope of the initiative is to define a model that is agnostic to cancer type and independent from specific technical standards, providing a common conceptual and logical backbone that can be progressively extended over time.

The scope of the Implementation Guide (<https://build.fhir.org/ig/hl7-eu/cancer-common/index.html>) is to document this common model and to support its implementation by providing mappings of the model to HL7 FHIR and to OMOP, enabling interoperability between primary data exchange and secondary data use in research contexts.

By defining consistent concepts, relationships, and longitudinal structures, the model enables—among other outcomes—the reconstruction of a typical cancer journey.

This project aims to define a minimal, extensible, and non-exhaustive European cancer data model that is:

- Agnostic to the type of cancer;
- Usable across different use cases;
- Leveraging experiences of European projects working with both primary and secondary use of data;
- Taking into account the availability and usability of reliable information in EHR systems.

The model is intended to be:

- Inclusive and cross-cutting, transversal to different cancer domains and purposes of use;
- Community-driven, considering the needs of different communities;
- Incremental, starting from a minimal core set and extending the model through subsequent iterations.

At the same time, ECCDM explicitly addresses the relationship between primary and secondary use of cancer data. Secondary use (e.g. research, analytics, AI) depends on the quality and structure of data captured during primary care. Understanding secondary-use requirements is therefore essential to identify which information should be collected during primary care in a consistent and reusable way.

ECCDM treats primary and secondary use as a single continuum: secondary-use needs inform the definition of minimum requirements for primary data capture, enabling reliable reuse without redefining data semantics at later stages.

ECCDM builds on experience gained across multiple European and national initiatives addressing cancer data.

Inputs include several European projects (e.g. IDEA4RC, PanCareSurPass, FLUTE etc.), as well as shared datasets and initiatives such as MEDOC (Minimal Essential Description of Cancer). Complementary national initiatives, such as OSIRIS, have also contributed relevant experience and requirements.

ECCDM is informed by existing international specifications, including the HL7 US mCODE Implementation Guide, which provides valuable implementation experience in oncology. ECCDM, however, addresses European interoperability needs through a model-driven approach based on conceptual and logical foundations.

In the European context, OMOP is widely adopted for research and secondary use, while HL7 FHIR is increasingly used for data exchange and interoperability. ECCDM is intended to provide a shared foundation that supports both representations and enables systematic mapping between OMOP and FHIR.

Example scenarios supported by ECCDM include:

- A multidisciplinary cancer team retrieves cancer-related data to assess the current diagnosis;
- Cancer data is stored in a research warehouse for use in various research projects;
- Representing cancer-related information consistently across projects and care settings;
- Enabling structured primary data capture that supports downstream research and secondary reuse;
- Supporting alignment and mapping between HL7 FHIR-based data exchange and OMOP-based research infrastructures.

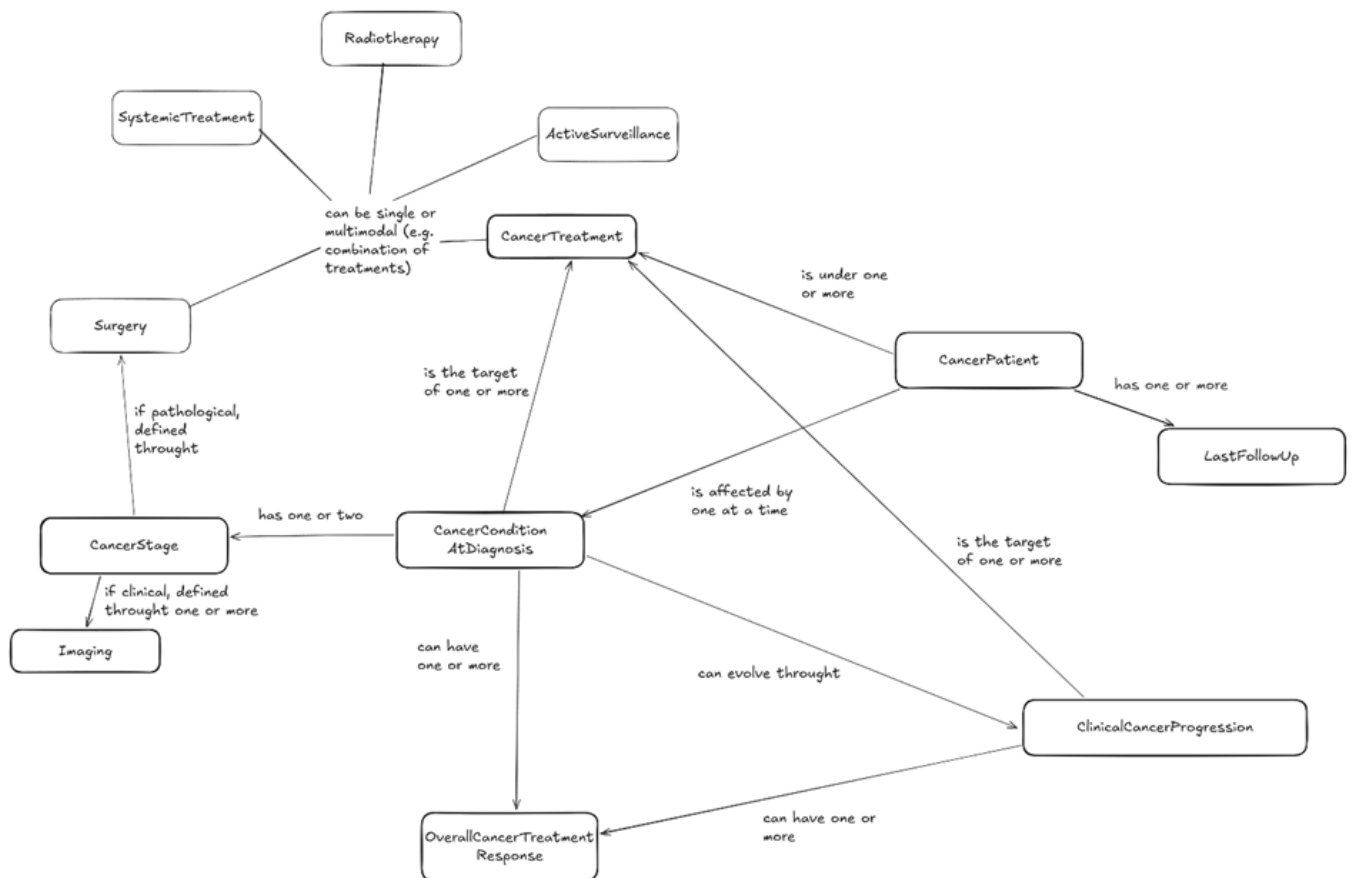
HL7 Europe has contributed in the context of the FLUTE project to the creation of the conceptual model

At a high level, the European Cancer Common Conceptual Model represents how a Cancer Patient is affected by one or more Cancer Conditions, how each condition is characterised at diagnosis, staged, treated, evaluated, and how it evolves over time.

The model captures the core structure of a typical cancer journey, including:

- the patient as the subject of care;
- the diagnosis of one or more cancer conditions;
- staging assessments associated with each condition;
- treatments administered with an explicit intent;
- evaluation of treatment response based on clinical evidence;
- clinical progression of the disease over time;
- follow up information summarising the most recent known patient status.

This overview introduces the main building blocks of the model before presenting the complete conceptual representation.



**Figure 2. European Cancer Common Conceptual Model**

As well as the definition of the patient cancer journey model.

The journey usually begins with one or more sources of clinical evidence (e.g. imaging examinations, laboratory tests, biopsies). These evidences support the assertion of a cancer diagnosis. In some cases, additional evidence may be required to further characterise the cancer condition, for example to refine cancer stage or tumour characteristics (e.g. tumour grade).

Once a cancer condition has been asserted, clinicians may decide to initiate a treatment. Treatments can occur in one or more episodes and may be of different types (i.e. surgery, radiotherapy, systemic treatment, active surveillance).

After a treatment has been delivered, new clinical evidence is collected. This post treatment evidence is required to support a follow up visit, during which the effect of the treatment is evaluated, the current status of the disease is determined and the next steps are defined. This cycle of evidence collection → assessment → decision may repeat multiple times over the course of the disease.

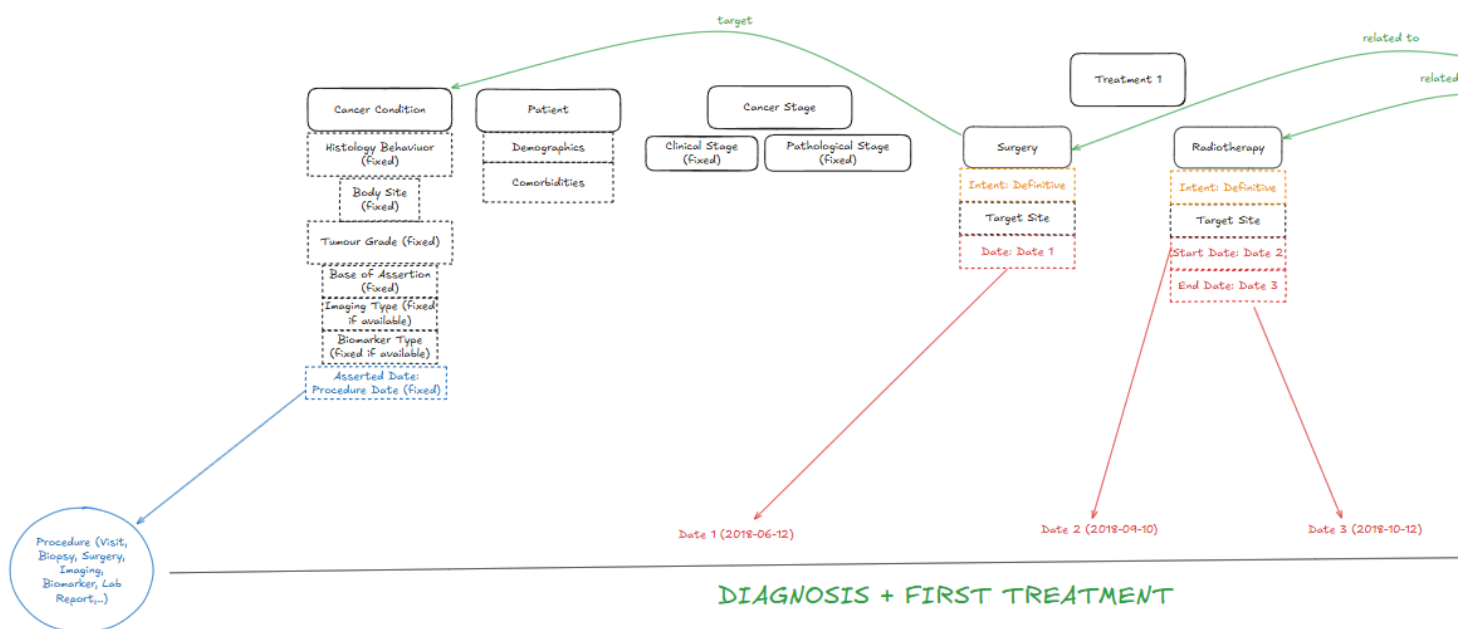


Figure 3. Extract of the European cancer patient typical Journey

### 1.2.3 Collaboration with Codex

To ensure alignment between European and US cancer (oncology) interoperability efforts, a dedicated collaboration between the HL7 Codex Accelerator and the PHOENIX initiative was initiated. This direction was agreed during the Common Cancer Data Model track held in Cologne, where stakeholders recognised the importance of avoiding divergent modelling approaches and fostering convergence across regions.

The collaboration aims to explore practical ways for Codex and PHOENIX to work together, including coordination of modelling priorities, comparison of artefacts, and structured exchange of implementation feedback. A key ambition of this joint effort is to assess the feasibility of defining, in the longer term, an international cancer model that would leverage and reconcile both the US Codex-driven approach and the European PHOENIX-driven modelling perspective.

Through this collaboration, HL7 Europe contributes to establishing a cross-regional pathway towards reusable and globally aligned cancer data standards, supporting both European needs and broader interoperability objectives.

### 1.3 Defining shareable research artefacts

One of the topics explored within FLUTE was the structuring of a consistent representation of research studies as shareable research artefacts. This work aimed to support reusability and interoperability across organisations and platforms, by enabling studies and associated resources to be described in a standardised and exchangeable manner.

To contribute to this objective, HL7 Europe engaged with the relevant HL7 FHIR work group(s) during HL7 Working Group Meetings (WGMs). These discussions focused on identifying appropriate modelling approaches and existing HL7 FHIR building blocks that could support the representation of study-related artefacts in a way that is compatible with federated research and cross-border collaboration requirements.

Through these interactions, HL7 Europe ensured that FLUTE reflections and needs were connected to ongoing HL7 standardisation activities, and that emerging HL7 guidance could be considered to strengthen the sustainability and shareability of research artefacts beyond the project context.

## Collaborations with other Initiatives

FLUTE, with its focus on federated research, including medical research, aims to enable secure data sharing and automate key elements of the research lifecycle. This positioning naturally aligns the project with several complementary European and international initiatives working on clinical research interoperability, diversity, oncology data models, and digital trial enablement. Some are outlined below:

### **READI**

The READI (Research in Europe and Diversity Inclusion) project is a European initiative dedicated to improving inclusivity and representativeness in clinical studies. It addresses systemic barriers that limit participation of underserved and underrepresented populations in clinical trials across Europe.

READI develops: Tools and operational frameworks to improve recruitment and retention, training materials to enhance diversity awareness, governance models to support inclusive trial design and execution.

### **HL7 Vulcan, CDISC and Clinical Trials Interoperability**

In parallel, HL7 Europe contributes to aligning HL7 FHIR-based artefacts with related global standards initiatives, including, HL7 Vulcan projects and CDISC standards for clinical research data through projects and workstreams, such as the Utilizing the Digital Protocol (UDP) and Schedule of Activities (SoA). This alignment supports semantic consistency between federated research infrastructures like FLUTE and structured clinical trial execution frameworks.

HL7 Europe continuously monitors and engages in developments related to digital clinical trials, notably through HL7 Vulcan and CDISC. Vulcan works to bridge healthcare and research standards using HL7 FHIR, enabling interoperability between clinical care systems and research platforms. In parallel, CDISC curates global regulatory standards, such as the Study Data Tabulation Model (SDTM) and the Analysis Data Model (ADaM), which support submissions for medical product marketing authorisations in the United States and Japan. CDISC also maintains a set of robust oncology standards focused on solid tumors. Mappings for common clinical concepts are available from the xShare project to form a pathway from HL7 FHIR to CDISC regulatory standards. For FLUTE, key areas of relevance include digital protocol representation, structured modelling of eligibility criteria, and the definition of reusable research artefacts that enable cross-platform automation the possibility of bridging healthcare data and clinical research submission data. This alignment ensures that FLUTE's federated research approach remains consistent with evolving international best practices in clinical trial digitisation.

### **MCODE Common Oncology Data Elements**

HL7 Europe also monitors oncology-specific standardisation efforts such as mCODE (Minimal Common Oncology Data Elements), which defines a standardised core set of oncology data elements to be exchanged via HL7 FHIR.

Developed in collaboration with IHE and oncology stakeholders, mCODE promotes consistent representation of cancer diagnoses, staging, genomics, and treatment data within electronic health records.

The convergence between: mCODE data structures, PHOENIX European modelling efforts,

And Vulcan's clinical research alignment activities, creates a broader interoperability ecosystem within which FLUTE's cancer-focused federated learning and research artefacts can be positioned sustainably.

### **Overall Alignment**

Through structured collaboration and continuous monitoring of related initiatives, HL7 Europe ensures that FLUTE: Contributes to, and benefits from, international interoperability developments, avoids fragmentation across oncology and research data model and anchors its federated research outputs within durable standardisation trajectories. This ecosystem-based alignment reinforces FLUTE's long-term sustainability and strengthens its relevance within the emerging EHDS landscape.

## **1.4 Link with DICOM imaging**

Imaging is identified as one of the six priority domains of the EHDS, reflecting its central role in both clinical care and secondary use scenarios such as research and AI development. Ensuring interoperability between imaging data and clinical data is therefore a key requirement at the European level.

In this context, HL7 Europe has been actively contributing to the definition of a standardised format for imaging reports based on FHIR, with a strong likelihood of adoption within the EHDS framework. This work is formalised through the HL7 Europe Imaging Report Implementation Guide (R5), which defines how imaging-related clinical information can be structured, exchanged, and reused consistently across primary and secondary care contexts.

(Reference: <https://build.fhir.org/ig/hl7-eu/imaging-r5/en/index.html>)

This model is designed to support both routine clinical workflows (primary use) and secondary use cases such as research, analytics, and AI, ensuring continuity and consistency of data across the care continuum.

Within the FLUTE project, imaging data was handled through a complementary approach: images were pseudonymised and stored in a DICOM repository, while associated metadata needed to be made available within a FHIR-based environment to enable integration with study workflows, cohorting, and data extraction processes.

To address this, HL7 Europe worked in parallel on two complementary tracks:

- On one hand, defining and structuring imaging metadata within the FLUTE sandbox, including mechanisms to capture such metadata through structured questionnaires and integrate it into the FHIR repository,
- On the other hand, contributing to the broader European standardisation effort for imaging reports within HL7 Europe, aligned with EHDS objectives.

Due to project time constraints, full convergence between these two tracks could not be achieved within the FLUTE project timeline. However, the mechanisms implemented in the sandbox — in particular the structured capture and integration of imaging metadata into FHIR — are compatible with the direction of the HL7 Europe imaging report profiles.

As a result, there is a strong expectation that future iterations of the FLUTE platform will be able to align with, and potentially adopt, the HL7 Europe Imaging Report profiles, enabling full interoperability between DICOM imaging repositories and FHIR-based research and care workflows.

Overall, this work positions FLUTE within the broader European trajectory for imaging interoperability, ensuring that project outcomes remain compatible with emerging EHDS standards and can evolve towards full compliance in future developments.

## 1.5 Standardisation of AI integration

The integration of Artificial Intelligence into healthcare interoperability frameworks is an emerging topic within HL7 (<https://info.hl7.org/ai>). In this context, HL7 Europe established a liaison with the HL7 International AI Task Force and actively participated in the AI track during the HL7 International Working Group Meeting in Pittsburgh.

Within FLUTE, HL7 Europe explored how existing FHIR standards could support key aspects of the AI lifecycle, with a particular focus on training data preparation and deployment of trained models in clinical workflows.

A first area of work consisted in structuring the knowledge of study variables required for AI training. These variables — defined within the context of a research study — were modelled using FHIR artefacts (notably EvidenceVariable, ResearchStudy, and associated Library resources) and evaluated on patient data to generate consistent and reproducible datasets. Mechanisms were also implemented to export these evaluated variables into tabular formats (e.g. CSV), making them directly usable for AI training pipelines.

A second area of exploration focused on the integration of trained AI models into care workflows. HL7 Europe investigated how previously defined FHIR artefacts representing

study variables could be reused as inputs for AI inference, ensuring consistency between training and deployment phases. In this context, CDS Hooks was used as a standard mechanism to integrate AI-based decision support into clinical systems, enabling real-time evaluation based on structured inputs derived from FHIR resources.

This approach demonstrates how existing HL7 standards can support both the preparation of AI-ready datasets and the integration of AI models into clinical workflows, without introducing new proprietary formats. It also contributes to ongoing HL7 reflections on AI transparency, reproducibility, and interoperability, by promoting the reuse of shared artefacts across the full AI lifecycle.

While the standardization and experimental work realized could not be put into practice in this version of the FLUTE platform due to time constraints, the result can be reused in future versions and might help into creating a standardized European framework for AI training and integration into Clinical Decision Support systems.

## 2 HL7 contribution to partners

### 2.1 Training

As part of its capacity-building activities, HL7 Europe contributed to FLUTE training efforts conducted in collaboration with the XiA project. These sessions focused on strengthening partners' understanding and practical use of HL7 FHIR, interoperability principles, and standards-based modelling approaches relevant to federated research and secondary use of health data. This collaboration enabled knowledge exchange across projects while maintaining clear separation of resources and funding streams.

This coordinated approach enhanced technical capacity across both initiatives, promoted consistency in standards adoption, and fostered a shared understanding of interoperable research infrastructures within the broader European digital health ecosystem.

List trainings developed and delivered:

- EHDS data domains
- FHIR
- IGs
- FSH
- SUSHI MenuXML
- SUSHI Narrative and pagination
- SUSHI
- the Problematic of Medical Terminology
- Medical Terminologies
- Terminology Server
- ConceptMap
- Logical Model
- StructureMap
- Building a FHIR Mapping
- CQL
- CDS Hooks
- Evidence-Variable

Specific training developed for FLUTE specifically included:

- FHIR Questionnaire
- FHIR SDC Questionnaire
- Mapping Data to FHIR Questionnaire

### 2.2 Specification designing and sharing

HL7 Europe contributed to the FLUTE project by designing and publishing two FHIR Implementation Guides, one based on FHIR R4 and the other on FHIR R5:

<https://build.fhir.org/ig/hl7-eu/flute/>

<https://build.fhir.org/ig/hl7-eu/flute-requirements/>

These Implementation Guides consolidate the artefacts developed and explored throughout the project, covering the different concepts required to support FLUTE use case, including data acquisition, cohort definition, variable evaluation, and data extraction.

They serve a dual purpose:

- Documenting the modelling approaches and standards-based mechanisms explored during the project,
- Providing the concrete artefacts used in the FLUTE experimentation environment.

In particular, the Implementation Guides include the FHIR artefacts (e.g. Questionnaire, StructureMap, ResearchStudy, EvidenceVariable, Library) that are actively used to generate structured datasets within the project. These datasets are then leveraged for research workflows and AI training.

By publishing these specifications openly, HL7 Europe ensures that the work performed within FLUTE is reusable, transparent, and aligned with broader HL7 and European interoperability initiatives, enabling other projects to build upon the same standards-based foundations.

### 2.3 Workshop with partners

Throughout the project, HL7 Europe organised and contributed to a series of technical workshops with FLUTE partners, focusing on enabling the practical implementation of data acquisition, interoperability, and research workflows.

These workshops were conducted with a pragmatic objective: to support partners in defining and producing usable project artefacts, ensuring that the experimentation could be carried through to completion in a consistent and standards-based manner.

Key topics addressed during these workshops included:

- The definition of a structured data acquisition workflow based on FHIR questionnaires,
- The design and refinement of Questionnaire artefacts adapted to FLUTE study requirements,
- The implementation of extraction mechanisms to transform collected data into structured FHIR resources,
- The alignment between clinical data capture and imaging workflows, including the integration of DICOM metadata into FHIR,

- The validation and iterative improvement of questionnaire-based workflows through partner feedback.

Through this collaborative work, HL7 Europe supported partners in operationalising a complete pipeline from data capture to data exploitation. The artefacts produced, in particular the questionnaires and associated extraction logic, are actively used within the project to generate structured datasets, which are then leveraged for AI training and research purposes.

This hands-on approach ensured that standards were not only defined but effectively applied within the project, enabling a functional and reproducible experimentation environment.

## 2.4 Standard usage illustration

Within FLUTE, HL7 Europe was tasked by the consortium to explore how existing interoperability standards could support federated research mechanisms, going beyond traditional cancer data modelling activities. Federated research, and particularly the use of computable study definitions, distributed cohort generation, and standard-based variable evaluation, remains an emerging domain where formalised use of interoperability standards is still limited.

In response to this challenge, HL7 Europe investigated how established HL7 FHIR artefacts could be reused and combined to meet the project's requirements without introducing proprietary mechanisms. Rather than developing new standards, the approach focused on leveraging existing resources such as **ResearchStudy**, **EvidenceVariable**, **Questionnaire**, **CQL Library**, and standard FHIR operations to operationalise study definitions, eligibility criteria, and variable computation in a federated context.

The following subsections illustrate how these standards were applied in practice to support distributed cohort identification, structured datamart generation, and reuse of harmonised variables for AI-enabled clinical workflows. This demonstrates that existing HL7 standards and others, when appropriately combined, can support advanced federated research scenarios aligned with the project's objectives.

### 2.4.1 Data acquisition

#### User scenario

A prostate cancer study has been defined, and its structured questionnaire has been shared with participating hospitals in advance.

During a routine consultation, a urologist identifies a patient who may be relevant for the study.

The practitioner opens the predefined study questionnaire within the hospital system. The questionnaire contains the agreed clinical variables required for the research protocol (e.g., PSA value, PIRADS score, prostate volume, biopsy type, family history).

The practitioner reviews the patient's unstructured medical record and enters the required information into the questionnaire.

Once completed, the questionnaire is saved in the hospital system and the data from the questionnaire feeding the patient medical record. The practitioner also gets feedback from clinical decision support systems based on the study inclusion criterias which tells him whether the patient is eligible for the study.

The patient's structured study data is now available within the hospital environment for potential use according to the applicable governance procedures.

## Standards illustration

### Design and using the questionnaire

Once the researcher and doctors have identified the variables needed for training the AI models, a FHIR representation of those variables is made available to discussion to validate the concepts and the semantics behind in the form of an HL7 FHIR logical model. Those concepts are then mapped to a FHIR Questionnaire resource through a ConceptMap, allowing to make the data available in a QuestionnaireResponse format when instantiated and filled with data. It can then be checked on the Questionnaire for validity.

#### 5.3 Cross-border prostate cancer pilot study Requirements

Requirements for the AI models FLUTE should consider impacts on the use of developed AI models, in particular with respect to the impact of this art should be part of a broader requirements should be co-appropriate - align with the scenarios for AI development that are set in the requirements for AI in more detail

#### HL7 FHIR Logical Model

This structure is derived from [Base](#)

Name	Flags	Card.	Base	Description & Comments
studyVariables		0..*	Base	Prostate Cancer Instances of
patient		1..1	Reference(Patient)	Patient
ageAtBiopsy		1..1	integer	Age at the time
pcaFamilyHistory		1..1	integer	Family history
typeOfBiopsy		1..1	integer	Type of biopsy
psa		1..1	Quantity	Level of prostate
dre		1..1	integer	Results of digital
prostateVolume		1..1	Quantity	Prostate volume
pirads		1..1	integer	PI-RADS score

#### HL7 FHIR Questionnaire

Questionnaire: Study variable Extraction (Experimental)

Official URL: <http://hl7.org/fhir/ig/flute/questionnaire/questionnaire-studyvariable>

Active as of 2024-03-26 | Realm: **EU**

ItemID	Text	Cardinality	Type	Flags
StudyVariableExtraction	Study variable extractable questionnaire		Questionnaire	
procedure-group	null	0..1	group	
type-of-biopsy	Type of biopsy	1..1	choice	
procedure-status	Status			
age-at-biopsy	Age at biopsy			
biopsy	Biopsy	1..1	choice	
biopsy-bodySite	Biopsy bodySite		choice	
family-history-group	null		group	

#### Model Map

ConceptMap: Study variables Model to FHIR R4 Map (Experimental)

Official URL: <http://hl7.org/fhir/ig/flute/conceptmap/cn-studyvariablebase2024>

Draft as of 2023-12-08 | Realm: **EU**

Mapping from Prostate Cancer study Variables to http://hl7.org/fhir/ig/flute/questionnaire/questionnaire-studyvariable

Study variables Model to this guide Map

Group	Source Code	Relationship	Target Code
Group 1	StudyVariables.prostateVolume	is equal to	Observation.valueQuantity
Group 2	StudyVariables.psa	is equal to	Observation.valueQuantity
Group 3	StudyVariables.dre	is related to	Procedure.outcome
Group 4	StudyVariables.pirads	is equal to	Observation.valueQuantity
Group 5	StudyVariables.ageAtBiopsy	is equal to	Observation.valueQuantity

Questionnaire to FHIR Resources

Sample QuestionnaireResponse: Questionnaire-StudyVariable-S1

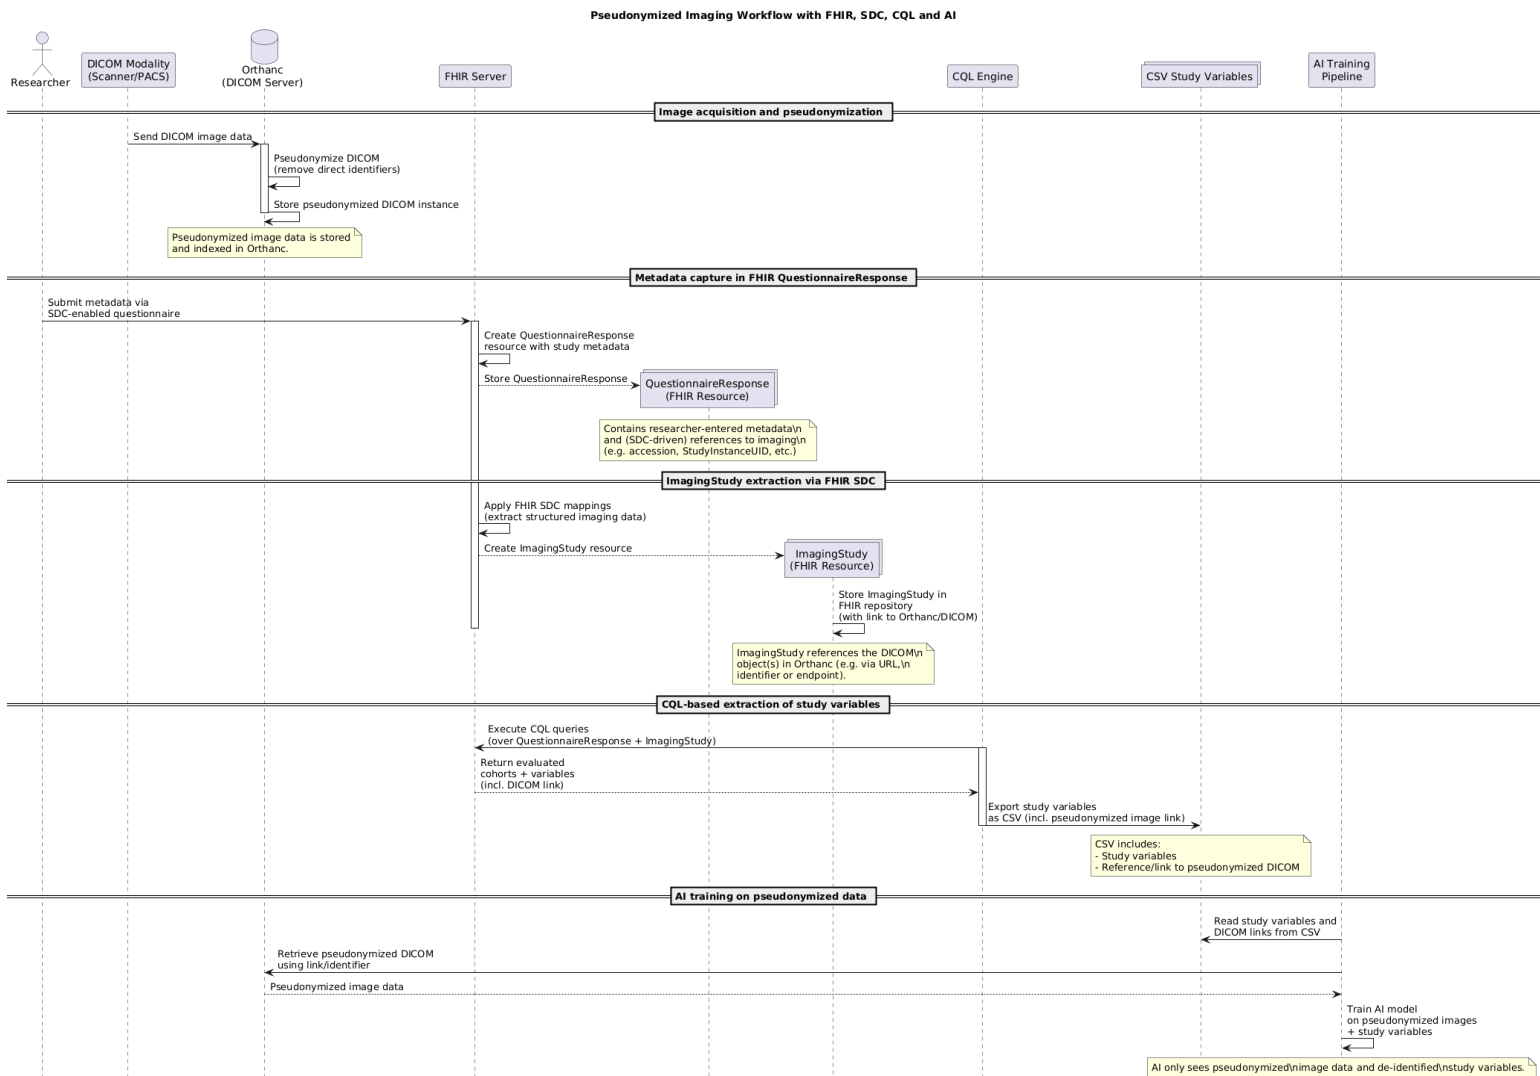
ItemID	Text	Answer	Definition
StudyVariableExtraction	Study variable extractable questionnaire	Questionnaire	Study variable extractable questionnaire
procedure-group	null		procedure-group
type-of-biopsy	Type of biopsy	SNOAED CT (all ve	SNOAED CT (all ve
procedure-status	Status	54 a	EventStatusCode
age-at-biopsy	Age at biopsy	54 a	SNOMED CT (all ve
biopsy	Biopsy	SNOAED CT (all ve	SNOAED CT (all ve
biopsy-bodySite	Biopsy bodySite	SNOAED CT (all ve	SNOAED CT (all ve
family-history-group	null		family-history-group
family-history-pca	Family history of prostate cancer	SNOAED CT (all ve	SNOAED CT (all ve
family-status	Family status	FamilyHistoryHist	FamilyHistoryHist
family-relationship	Family relationship	Relationship	Relationship
observation-psa-group	Observation group	Observation	Observation
psa	Measure of prostate-specific antigen reported in ng/ml	5 ng/mL	5 ng/mL

#### HL7 FHIR Questionnaire Response

Figure 4: Questionnaire response generation

## Making the link with DICOM imaging

Imaging data acquired in **DICOM** format is first pseudonymised and stored in the local imaging repository. Study-related imaging metadata is captured through a structured FHIR **Questionnaire**, generating a **QuestionnaireResponse** that includes references to the imaging identifiers (e.g., StudyInstanceUID). Using predefined mappings, an **ImagingStudy** resource is created in the FHIR repository and linked to the corresponding pseudonymised DICOM objects. During the **CQL**-based evaluation phase, both structured clinical variables and imaging-related references can be extracted.



**Figure 5: Pseudonymisation workflow**

## Checking eligibility

Once the **QuestionnaireResponse** is validated, it can be saved in a FHIR repository and extracted, meaning that FHIR resources are generated from the responses, defined in the **Questionnaire** definition. A clinical decision support system (**CDS-Hooks service**) is then triggered, which calls the evaluation of the inclusion criteria of the study using the **\$evaluate** operation on the FHIR **Library** resource. The software used then gets immediate visual feedback in the form of a **CDS-Hooks card** informing of the eligibility of the patient for the study. Managing Patient consent for the study can also be done in this step by enforcing part of the eligibility on presence or absence of a FHIR **Consent** resource.



Figure 6: Checking eligibility workflow

### 2.4.2 Make data available for research

#### User scenario

After structured questionnaires have been completed for multiple patients, the hospital activates the research protocol within its local FLUTE infrastructure.

The system automatically applies the predefined eligibility criteria to the structured patient data.

Patients who meet the inclusion conditions are identified without manual intervention.

For those patients:

- The relevant study variables are automatically extracted.
- The data is transformed into the predefined research format.
- The resulting dataset is stored in a secured local datamart dedicated to research activities.
- This datamart remains within the hospital's protected environment.

Once prepared, the structured dataset becomes available for federated research processes, allowing participation in multi-centre analysis while keeping patient-level data under institutional control.

## Standards illustration

### Evaluation of the variables for one patient

Once the **QuestionnaireResponse** is validated and stored in the FHIR repository, the study variables for that specific patient can be evaluated. The system invokes the **\$evaluate** operation on the corresponding FHIR **Library** resource, which contains the **CQL** logic defining how the study variables must be computed. The evaluation engine reads the structured patient data, applies the predefined rules, and generates the calculated variables in a standardised FHIR **Parameters** resource. Pseudonymization levels can also be described in the CQL logic, and applied at the evaluation, making sure that data is made imprecise on a study per study base.

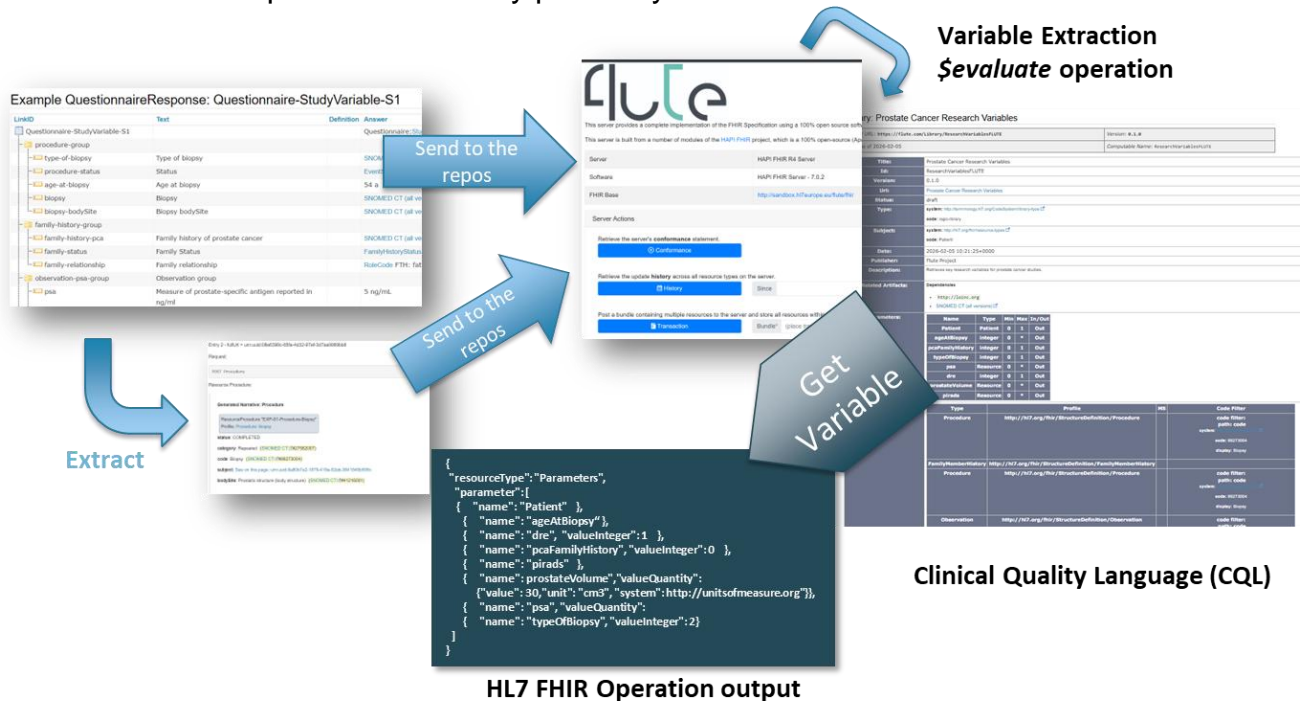


Figure 7: Evaluation of the variable for one patient workflow

## Making the data available in other formats

Once the study variables for a patient have been evaluated and returned as a FHIR **Parameters** resource, the data can be transformed into alternative formats required for research processing. Using the **HL7 FHIR Mapping Language** and the associated **StructureMap**, the structured output is mapped to a tabular representation, such as a CSV file. This transformation preserves the semantic definition of each variable while enabling compatibility with statistical analysis tools and federated learning frameworks. The conversion process remains traceable and standard-based, ensuring that the exported dataset corresponds exactly to the evaluated study variables defined in the research specification.

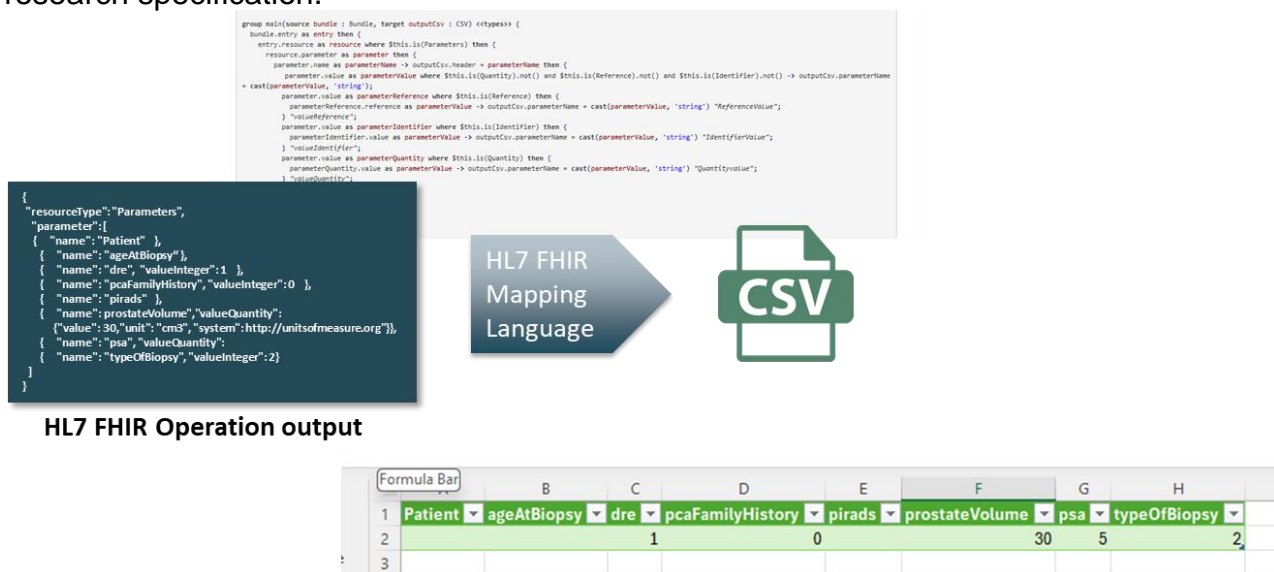


Figure 8: Making data available in multiple formats

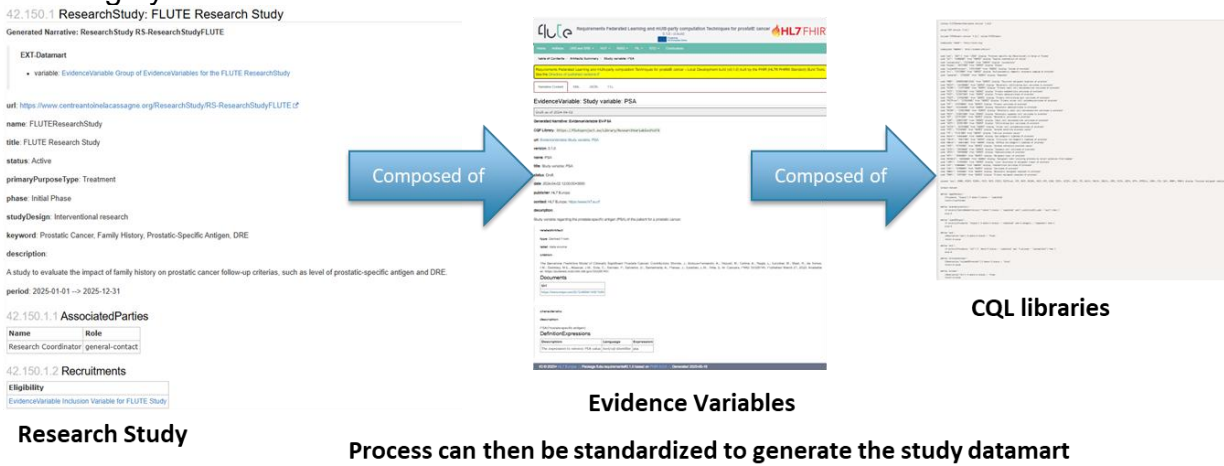
## Creation of a cohort and a datamart for a study

A study is first defined and shared as a standardised artefact using the **ResearchStudy** resource, which references the associated **EvidenceVariable** resources describing both the required study variables and the inclusion criteria as **CQL Library** resources. This study definition constitutes a shareable and computable research specification that can be distributed across participating institutions. “The implementation leveraged FHIR R5 resources to support study-level modelling features not fully available in R4.

Once deployed locally, the inclusion criteria defined in the study are applied to the entire FHIR repository of the institution. The system evaluates which patients satisfy the eligibility conditions based on the structured clinical data available.

For each patient meeting the criteria, the study variables defined under the **EvidenceVariable** components are then evaluated using the corresponding computation logic. The resulting harmonised variables are generated, converted and stored in a secured local datamart representing the cohort for that specific study.

Because the study definition, eligibility logic, and variable specifications are expressed as standardised artefacts, the same research protocol can be executed consistently across sites. Each institution independently generates its own cohort and datamart based on the identical shared specification, enabling federated research while preserving data sovereignty.



Patient	ageAtBiopsy	dre	pirads	prostateVolume	psa	typeOfBiopsy
e7d0a0fb1e53b6ced5fbed5796f4713f	59	0	3	2	40	0
277de635c18980e93dd118617d73ec26	43	0	3	2	20	0
4d7987ab3965ca521244bc4c0c606cd6	4	0	4	4	4	0
72780f69acc062f7558f4ff8084b66e8	4	1	3	4	8	0

Figure 9: Workflow - Creation of a cohort and a datamart for a study

### 2.4.3 Data acquisition – alternative post-AI training

#### User scenario

Following completion of the federated research phase, an AI model has been trained using the harmonised study data across participating institutions.

The trained model is deployed within the hospital’s clinical environment.

During a new consultation, a practitioner opens a structured clinical questionnaire similar to the one previously used for research purposes.

The practitioner enters the relevant patient information (e.g., PSA value, PIRADS score, biopsy details).

Once the questionnaire is completed, the system automatically evaluates the required study variables and sends only those computed variables to the locally deployed AI model.

Based on this input, the AI model generates a clinical output (for example, a risk score or decision support recommendation), which is returned to the practitioner within the consultation workflow.

Only the necessary evaluated variables are transmitted to the AI component, ensuring compliance with the principle of minimal data usage.

The practitioner receives the AI-supported insight as part of routine care and retains full responsibility for the final clinical decision.

## Standards illustration

### Making data available for AI evaluation in care

Once the study variables have been defined and their computation logic established, the same evaluation mechanism, and the same **CQL Library** to evaluate them can be reused in a clinical care context, displaying semantic continuity between research and care. When a practitioner completes the structured questionnaire for a patient, the system evaluates the required variables using the predefined logic and generates a harmonised set of computed values in a FHIR **Parameters** resource. These evaluated variables, rather than the full patient record, are then transmitted to the AI component through a **CDS-Hooks** call. The AI model receives only the minimal, standardised dataset necessary for inference and returns its output to the clinical system.

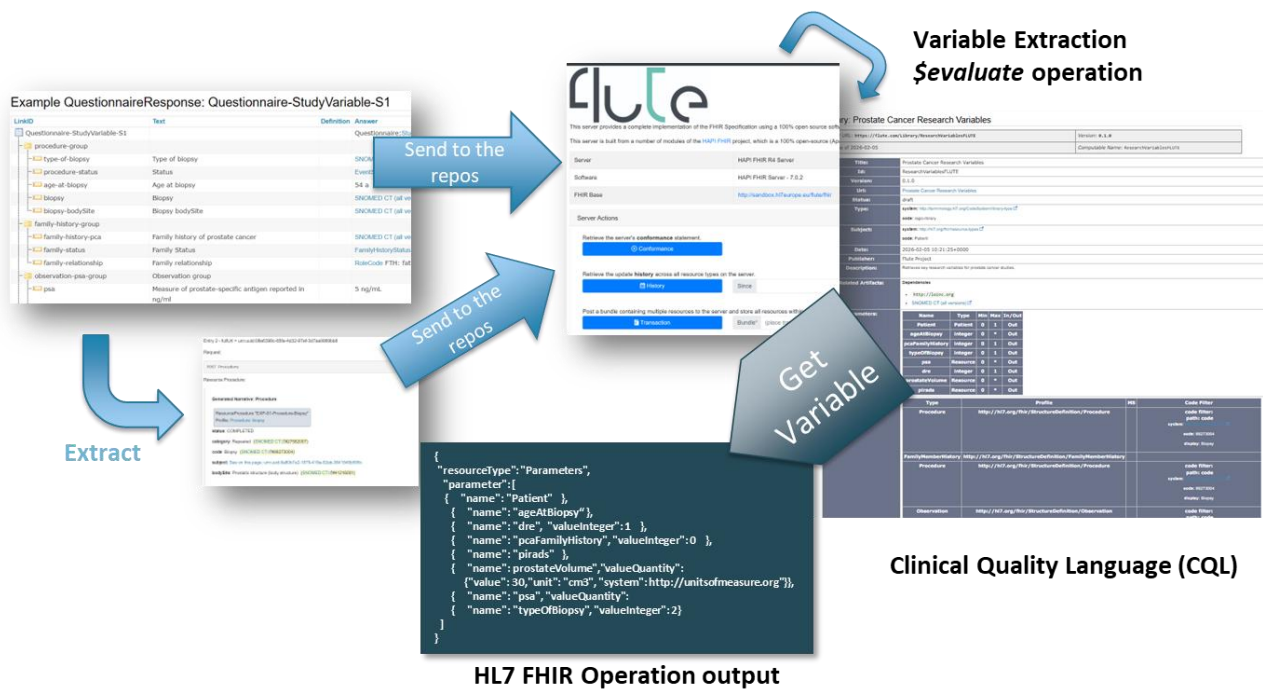


Figure 10: Workflow - Making data available for AI evaluation in care, reusing the evaluation of variables for a patient mechanism used in datamart generation

## 2.5 HL7 Europe Facilities - Sandbox

HL7 Europe maintains several supporting facilities that provide a secure, controlled environment for healthcare organizations, vendors, and developers to test interoperability solutions before live deployment. Through its initiatives, a sandbox enables validation of standards such as HL7 V2, CDA, and FHIR, supporting realistic data exchange scenarios. These facilities include conformance testing tools, validation services, terminology servers, security frameworks, training resources, and technical documentation. Together, they reduce implementation risk, strengthen compliance with European interoperability requirements, and foster collaboration across national programs, EU projects, and industry partners. Providing the FLUTE project with access to this structured testing environment can accelerate innovation while ensuring alignment with regulatory and cross-border data exchange objectives.

### 2.5.1 Architecture of the FLUTE experimental sandbox

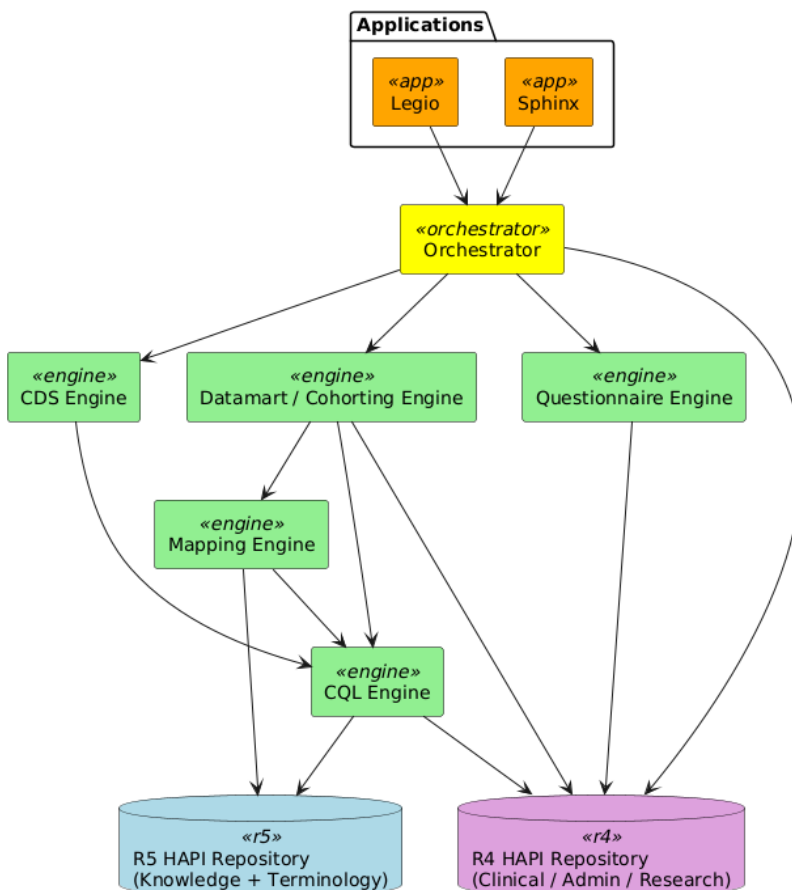


Figure 11: FLUTE experimental sandbox architecture

## 2.5.2 Open-source components of the architecture pre-existing to the project

### HAPI repository

The core of the sandbox architecture relies on a HAPI FHIR JPA Server Starter implementation. This open-source FHIR server provides a production-grade persistence layer and RESTful API compliant with HL7 FHIR specifications.

Within the FLUTE sandbox, both FHIR R4 and FHIR R5 versions were deployed in order to support different modelling requirements. R4 was used for compatibility with widely adopted clinical implementations, while R5 enabled advanced modelling features required for research artefacts such as enhanced ResearchStudy and EvidenceVariable resources.

The HAPI JPA Starter provides:

- Persistent storage of FHIR resources
- Versioning and history tracking
- Terminology support
- Search and filtering capabilities
- Some FHIR operations support

### Mirth connect

Mirth Connect is used within the FLUTE experimental sandbox as the orchestration layer coordinating interactions between the different architectural components. It acts as an integration engine responsible for routing, transforming, and triggering workflows across the platform.

Within the sandbox, Mirth Connect manages the exchange of messages between the FHIR repository, the CQL Engine, the Mapping Engine, the Cohorting Engine, the Datamart Engine, and external systems (such as imaging repositories or AI components). It enables controlled sequencing of operations, ensuring that data capture, evaluation, transformation, and export steps are executed in the correct order.

Concretely, it allows orchestration of multi-step workflows (e.g., Questionnaire submission → eligibility evaluation → variable computation → datamart generation).

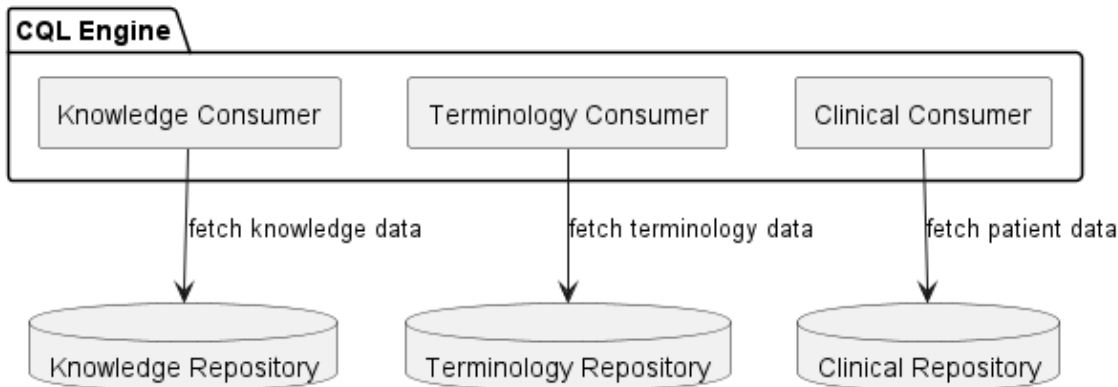
### CQL engine

The CQL Engine is an evaluation engine designed to interpret and execute expressions written in Clinical Quality Language (CQL), derived from the HAPI library to be deployed as a micro-service.

It is used to apply business rules, calculate indicators, and execute decision logic on FHIR data. It primarily relies on the Library and Measure resources and exposes standard operations such as `$evaluate` and `$evaluate-measure`, enabling the generation of actionable MeasureReport outputs.

In practical terms, it allows the execution of business rules on patient data, the calculation of quality or performance indicators, the production of structured results for business intelligence tools, and the support of clinical decision support mechanisms.

This component is leveraged whenever a project involves indicator computation, scoring, standardised statistical processing, or decision rules based on clinical data. It represents the core computational and logical evaluation capability within the platform foundation.



**Figure 12: CQL Engine diagram**

### Questionnaire engine

The Questionnaire Engine is an execution engine dedicated to FHIR questionnaires, compliant with the Structured Data Capture (SDC) Implementation Guide. It supports structured data collection workflows, including dynamic rendering, pre-population, and data extraction.

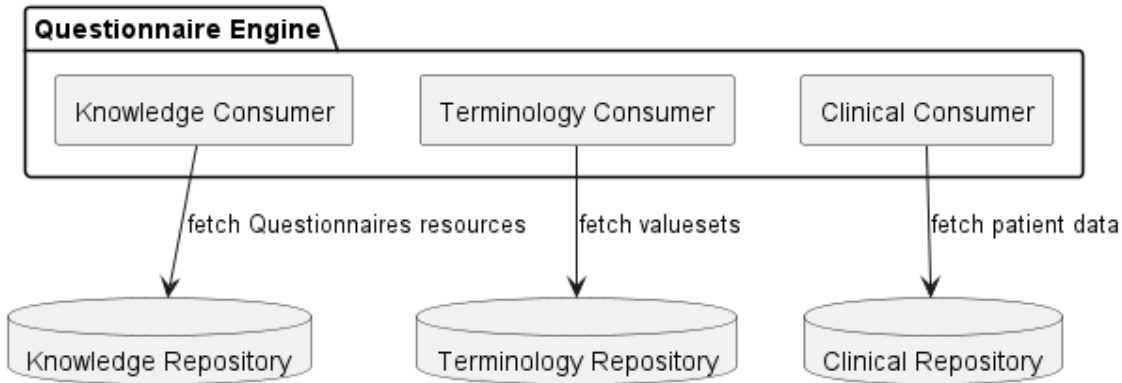
It exposes standard operations such as:

- `$populate`, which enables the pre-filling of a Questionnaire based on patient or contextual data,
- `$extract`, which transforms a QuestionnaireResponse into structured FHIR resources (for example, using StructureMap).

The engine relies on the Questionnaire, QuestionnaireResponse, StructureMap, and Parameters resources, and can interact with terminology services and existing clinical data repositories.

This component is leveraged when a project involves:

- the implementation of structured forms,
- the standardised collection of clinical or administrative data,
- automated pre-population from the patient record,
- transformation of responses into computable FHIR resources.



**Figure 13: Questionnaire engine diagram**

### Open-source GUI supporting CQL artefact development: Polus

In the context of modelling and validating computable logic within FLUTE, HL7 Europe relied in part on open-source tooling designed to facilitate the authoring and testing of Clinical Quality Language (CQL) artefacts.

One such tool is Polus, an open-source utility dedicated to the design, analysis, and testing of CQL-based FHIR artefacts. Polus supports the creation and validation of Library resources and enables structured testing of CQL expressions used in conjunction with FHIR resources such as Measure, PlanDefinition, or other decision-support mechanisms.

Within FLUTE-related activities, this tooling was used to support the exploration and validation of computable logic expressed in CQL, particularly in scenarios involving eligibility criteria evaluation, indicator computation, and study variable definition. Its purpose is to make CQL logic more transparent, testable, and maintainable within a FHIR-based architecture.

## 2.5.3 Components of the architecture developed or updated in the context of the project

### Questionnaire frontend component

To support the practical use of FHIR questionnaires within the FLUTE sandbox, HL7 Europe developed an open-source front-end component enabling the rendering and completion of FHIR Questionnaire resources.

This component, implemented using React, allows users to display structured questionnaires and capture responses in a user-friendly interface, while remaining fully compliant with the underlying FHIR data model. It supports the execution of questionnaire-driven workflows as defined in the project, including structured data capture aligned with study variables.

Within FLUTE, this component was used to operationalise the data acquisition process, enabling practitioners to input study-related information through standardised forms. The collected data is then processed and transformed into structured FHIR resources, feeding downstream processes such as variable evaluation, cohorting, and datamart generation.

The component is released as open-source software and is available through the HL7 Europe open-source space. It is intended as a reusable building block to facilitate the adoption of FHIR-based structured data capture in similar contexts.

Study variable Extraction

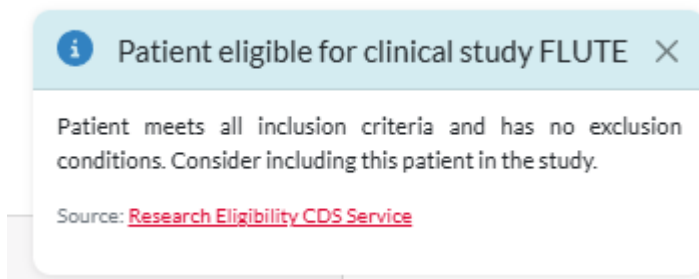
The screenshot displays two distinct questionnaire sections. The first section, titled 'Procedure Biopsy group', includes fields for 'Type of biopsy', 'Status' (set to 'completed'), 'Date when the biopsy was performed' (with a date picker), 'Biopsy', 'Biopsy bodySite' (set to 'Prostatic structure (body structure)'), and 'Family history of prostate cancer' (set to 'FAMMEMB'). The second section, titled 'Family history group', includes fields for 'Family history of prostate cancer', 'Family Status' (set to 'completed'), and 'Family relationship' (set to 'FAMMEMB').

**Figure 14. Example of visualization of the FLUTE questionnaire on the sandbox**

## CDS hooks card frontend component

To illustrate the integration of Clinical Decision Support within the FLUTE sandbox, HL7 Europe developed a front-end component implementing the CDS Hooks “card” concept.

This component enables the visualisation of decision support responses triggered during clinical workflows, based on standard CDS Hooks interactions. It displays contextual recommendations returned by a CDS service, such as patient eligibility for a clinical study.

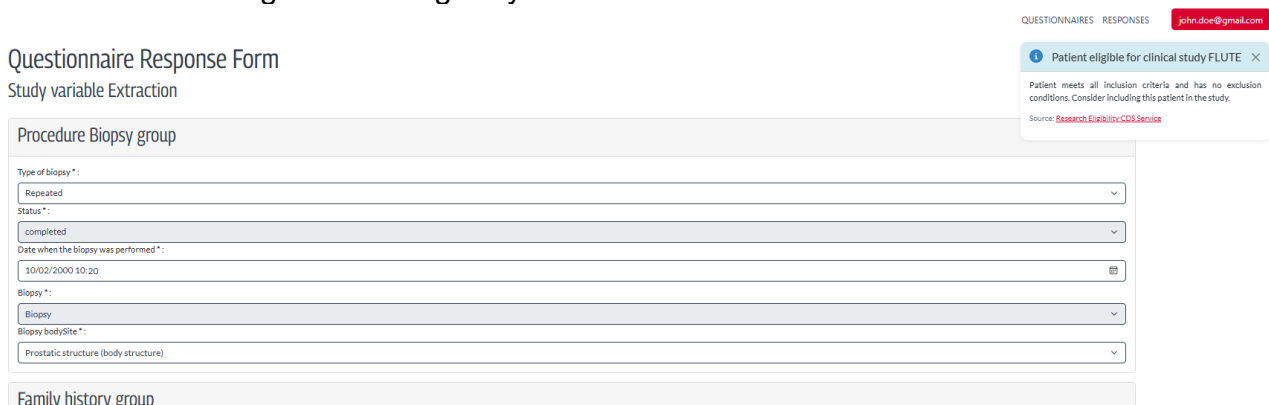


**Figure 15. Example of the CDS Hooks card component**

Within the FLUTE sandbox, this component is used in conjunction with questionnaire responses: once a QuestionnaireResponse is completed or reviewed, a CDS Hooks service is triggered to evaluate the study inclusion and exclusion criteria. The result is presented to the user through a CDS card, providing immediate feedback on whether the patient is eligible for the study.

This demonstrates how standardised decision support mechanisms can be integrated into data capture workflows, using FHIR-based artefacts and CDS Hooks without introducing custom interfaces.

The component is available within the sandbox environment and contributes to illustrating the end-to-end integration of eligibility evaluation into clinical workflows.



**Figure 16. Integration of the CDS hooks card component in an App**

## **Open-source GUI supporting Questionnaire & StructureMap artefact development: Janus (Contribution)**

To support the design and management of structured data capture and transformation artefacts within the FLUTE sandbox, HL7 Europe also relied on open-source tooling facilitating the creation of FHIR Questionnaire and StructureMap resources.

One such tool is Janus, an open-source application dedicated to the design and maintenance of FHIR artefacts used for structured data capture and data transformation. It provides an operational interface for producing and managing Questionnaire resources (supporting structured data collection) as well as StructureMap resources (supporting transformation and mapping workflows).

Within FLUTE-related activities, Janus was used to support the configuration and refinement of questionnaire definitions and transformation artefacts required for the experimental sandbox workflows. In this context, the tool was updated to integrate the developed questionnaire front-end component, enabling improved alignment between artefact design and runtime data capture mechanisms.

## **Open-source GUI supporting Questionnaire filling: Sphinx (Contribution)**

To facilitate the exploration and validation of structured data capture workflows within the FLUTE sandbox, HL7 Europe also relied on open-source tooling dedicated to the visualisation and operational use of FHIR questionnaires.

One such tool is Sphinx, an open-source application designed to support the exploration and utilisation of FHIR Questionnaire and QuestionnaireResponse resources. It provides a user-facing interface allowing teams to visualise questionnaire structures, instantiate responses, and review collected data in an operational, pedagogical, or research context.

Sphinx does not act as an execution engine or persistence layer; rather, it leverages the capabilities of the underlying FHIR platform to enable structured navigation and inspection of questionnaire artefacts and their responses. This makes it particularly suitable for design workshops, validation phases, demonstrations, quality monitoring, and secondary data usage scenarios.

Within FLUTE-related experimentation activities, Sphinx supported the validation and illustration of questionnaire-based workflows by making structured artefacts more accessible to both technical and domain experts. Its use helped bridge modelling artefacts and operational understanding without introducing additional architectural complexity.

## **Open-source GUI for study visualisation and datamart generation: Legio (Contribution)**

To support experimentation around cohort generation and research data preparation within the FLUTE sandbox, HL7 Europe also relied on open-source tooling designed to operationalise FHIR-based research workflows.

One such tool is Legio, an open-source application dedicated to the management and operational use of research artefacts within a FHIR-based environment. It enables the creation of cohorts, structuring of studies, and generation of protocol-specific datasets based on standardised criteria.

Legio supports the selection of patient populations according to defined inclusion and exclusion criteria, the extraction of study-relevant variables, and the preparation of datasets aligned with research protocols. It incorporates mechanisms supporting anonymisation and consent-aware processing, contributing to responsible and controlled secondary use of health data.

Within the context of FLUTE-related experimentation, Legio was used to illustrate and validate standards-based cohorting and datamart generation workflows built on FHIR resources such as ResearchStudy, EvidenceVariable, Group, and Library. Its use supported the operational testing of federated research scenarios without introducing proprietary modelling approaches.

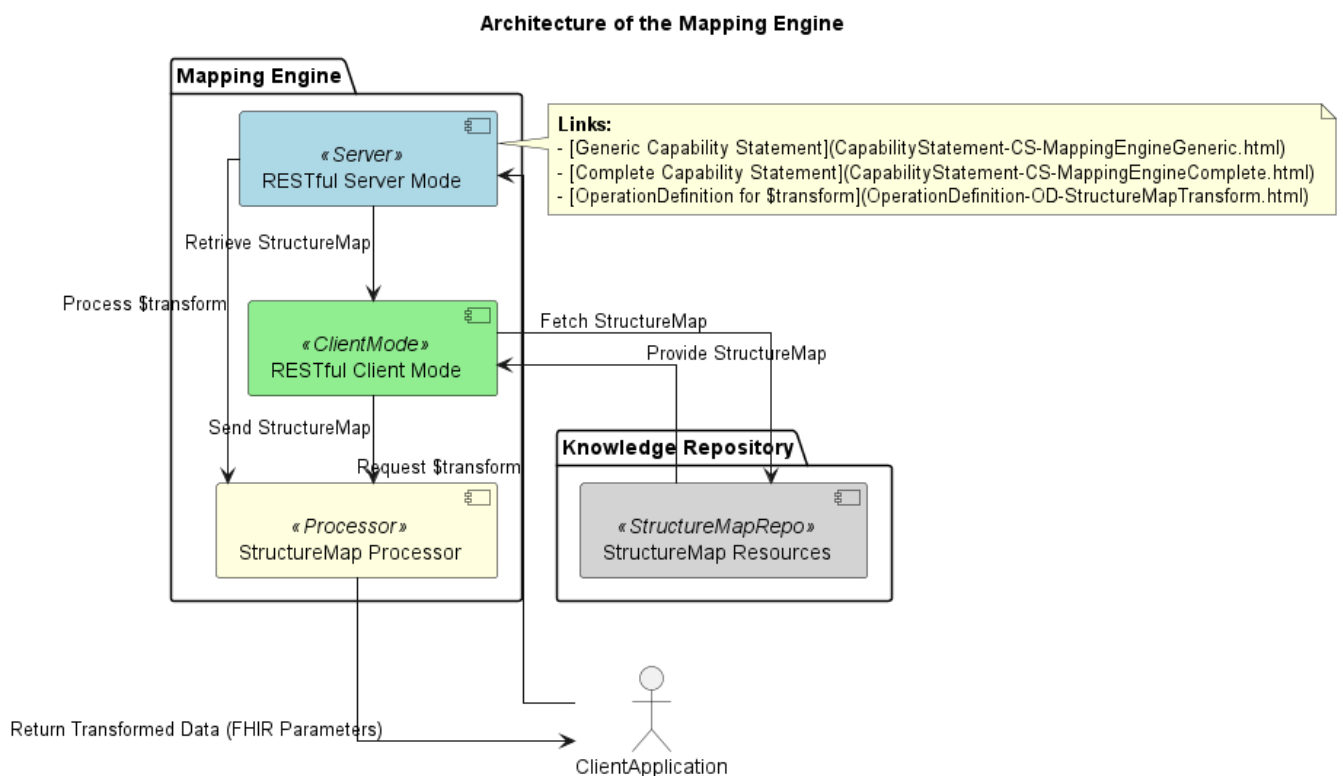
## Mapping engine (Evolution)

The Mapping Engine is a HAPI-based microservice designed to transform data from heterogeneous formats into FHIR resources, using StructureMap artefacts. It exposes the standard \$transform operation, enabling the application of a StructureMap to source content (JSON, XML, CSV, HL7v2, etc.) to produce a structured output.

In practical terms, it allows:

- Conversion of legacy data flows (HL7v2, CSV, JSON, etc.) to and from FHIR,
- Adaptation of data to a target profile,
- Industrialisation of reusable transformation rules,
- Execution of standardised transformation operations (create, copy, cast, translate, etc.).

This component is leveraged when a project involves data migration, integration with an existing information system, normalisation towards FHIR, or the implementation of an automated data ingestion pipeline.



**Figure 17: Architecture of the mapping engine**

In the context of the project, the mapping engine had to be evolved to support FHIR to csv conversions to generate the datamart.

## Cohorting engine (Creation)

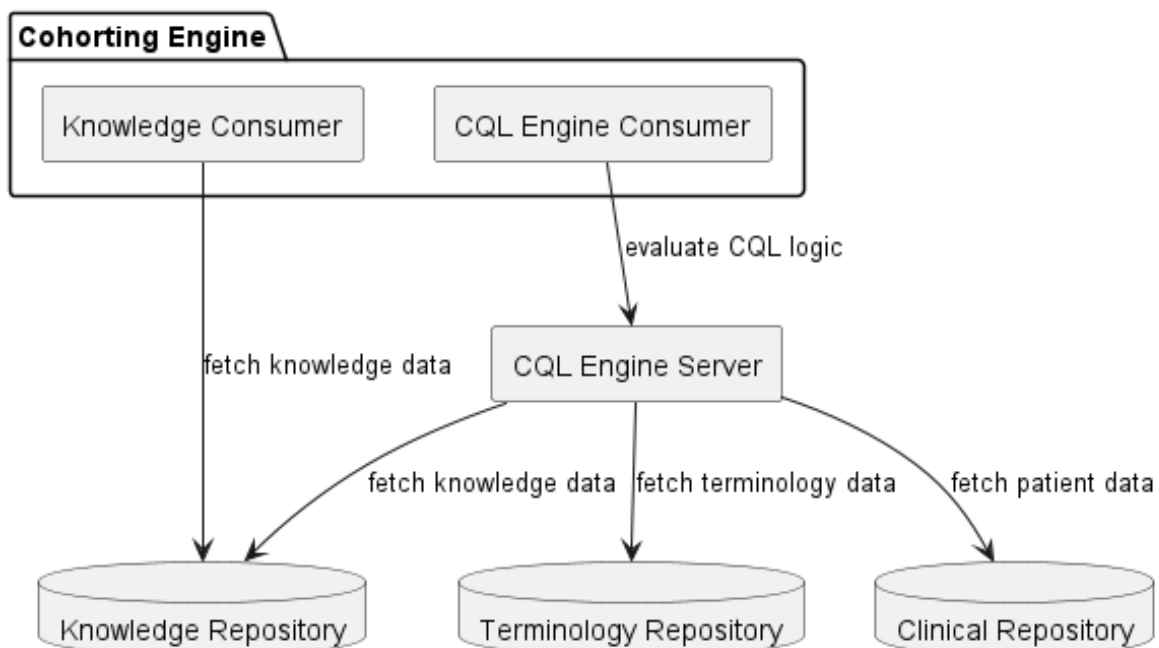
The Cohorting Engine is a HAPI-based microservice dedicated to patient selection and segmentation based on standardised clinical criteria. It enables the identification of eligible populations within the context of a study, quality monitoring initiative, or population-level analysis.

It relies in particular on the ResearchStudy, Group, EvidenceVariable, and Library resources, and can leverage CQL logic to define inclusion and exclusion criteria.

It exposes a cohorting operation that selects patients meeting the defined criteria and generates a Group resource representing the corresponding cohort.

This component is leveraged when a project involves:

- The creation of cohorts for a study or research protocol,
- Identification of patients based on specific clinical criteria,
- Recruitment for research studies,
- Population-level analyses based on formalised rule definitions.



**Figure 18: Cohort engine workflow**

This component was created in the context of the project as experimental development to test our standards usage hypothesis for creation of patient cohorts.

## Datamart engine (Creation)

The Datamart Engine is a HAPI-based microservice dedicated to the aggregation, transformation, and preparation of clinical data for analytical or research purposes. It enables the automated creation of structured datasets derived from FHIR resources.

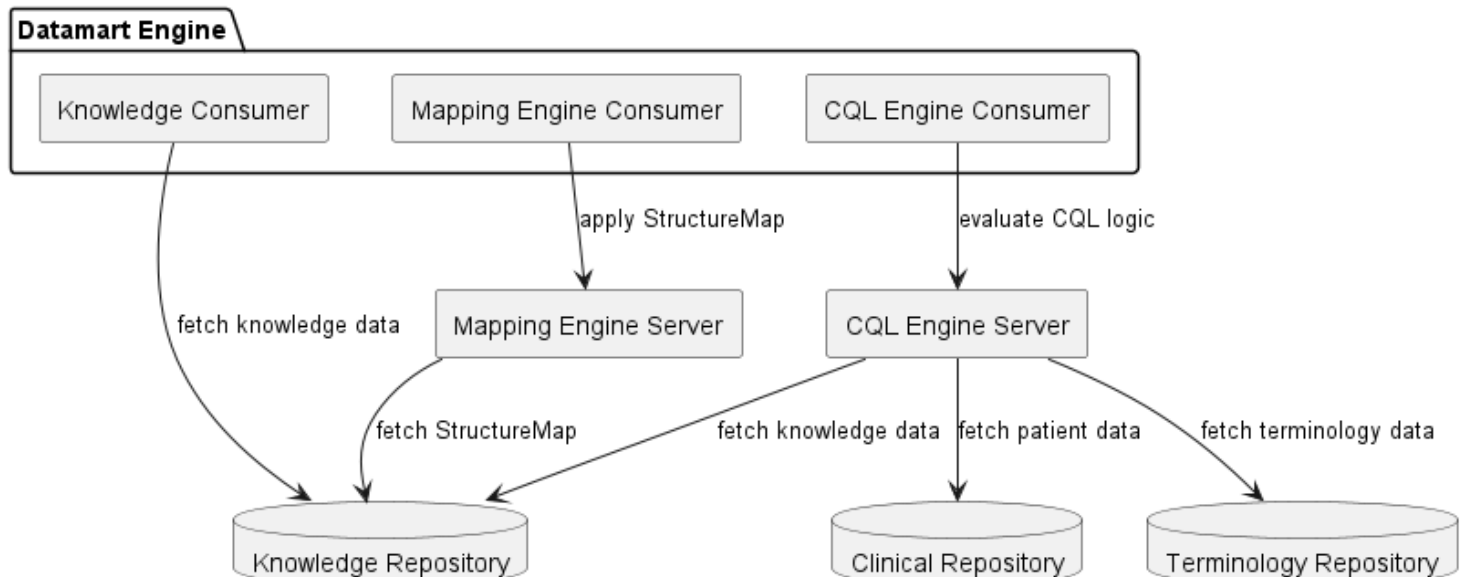
It exposes two main operations:

- \$generate-datamart, which aggregates and transforms data according to a defined context (for example, a ResearchStudy) and produces a structured set of evaluated variables,
- \$export-datamart, which allows the generated datamart to be exported in an actionable format (for example, CSV).

The engine can rely on study definitions, CQL libraries, and transformation rules (StructureMap) to generate consolidated datasets tailored for analysis.

This component is leveraged when a project involves:

- Preparing datasets for a research study,
- Producing secondary-use data from a FHIR-based foundation,
- Structured export to analytical or statistical tools,
- Establishing a reproducible research data pipeline.



**Figure 19: Datamart engine workflow**

This component was created in the context of the project as experimental development to test our standards usage hypothesis for generating a datamart for a specific study, especially in the context of federated research.

## **CDS-Hooks engine (Creation)**

The CDS Hooks Engine is a HAPI-based clinical decision support microservice based on the CDS Hooks specification and FHIR standards. It enables the integration of recommendations or alerts directly into clinical workflows (for example, during prescribing or consultation activities).

It operates by responding to events triggered by a client application (a hook), analysing the transmitted context, applying decision logic — potentially leveraging the CQL Engine — and returning structured recommendations in the form of “cards”.

In practical terms, this component allows the exposure of discoverable CDS services, real-time evaluation of business rules, and delivery of contextualised suggestions, information, or alerts within the clinical user interface.

The CDS Hooks Engine is leveraged when a project involves implementing an integrated decision-support mechanism within user workflows, whether for clinical, organisational, or quality-related use cases. It represents the platform’s cross-cutting real-time decision capability.

Architecture of the CDS Hooks Engine

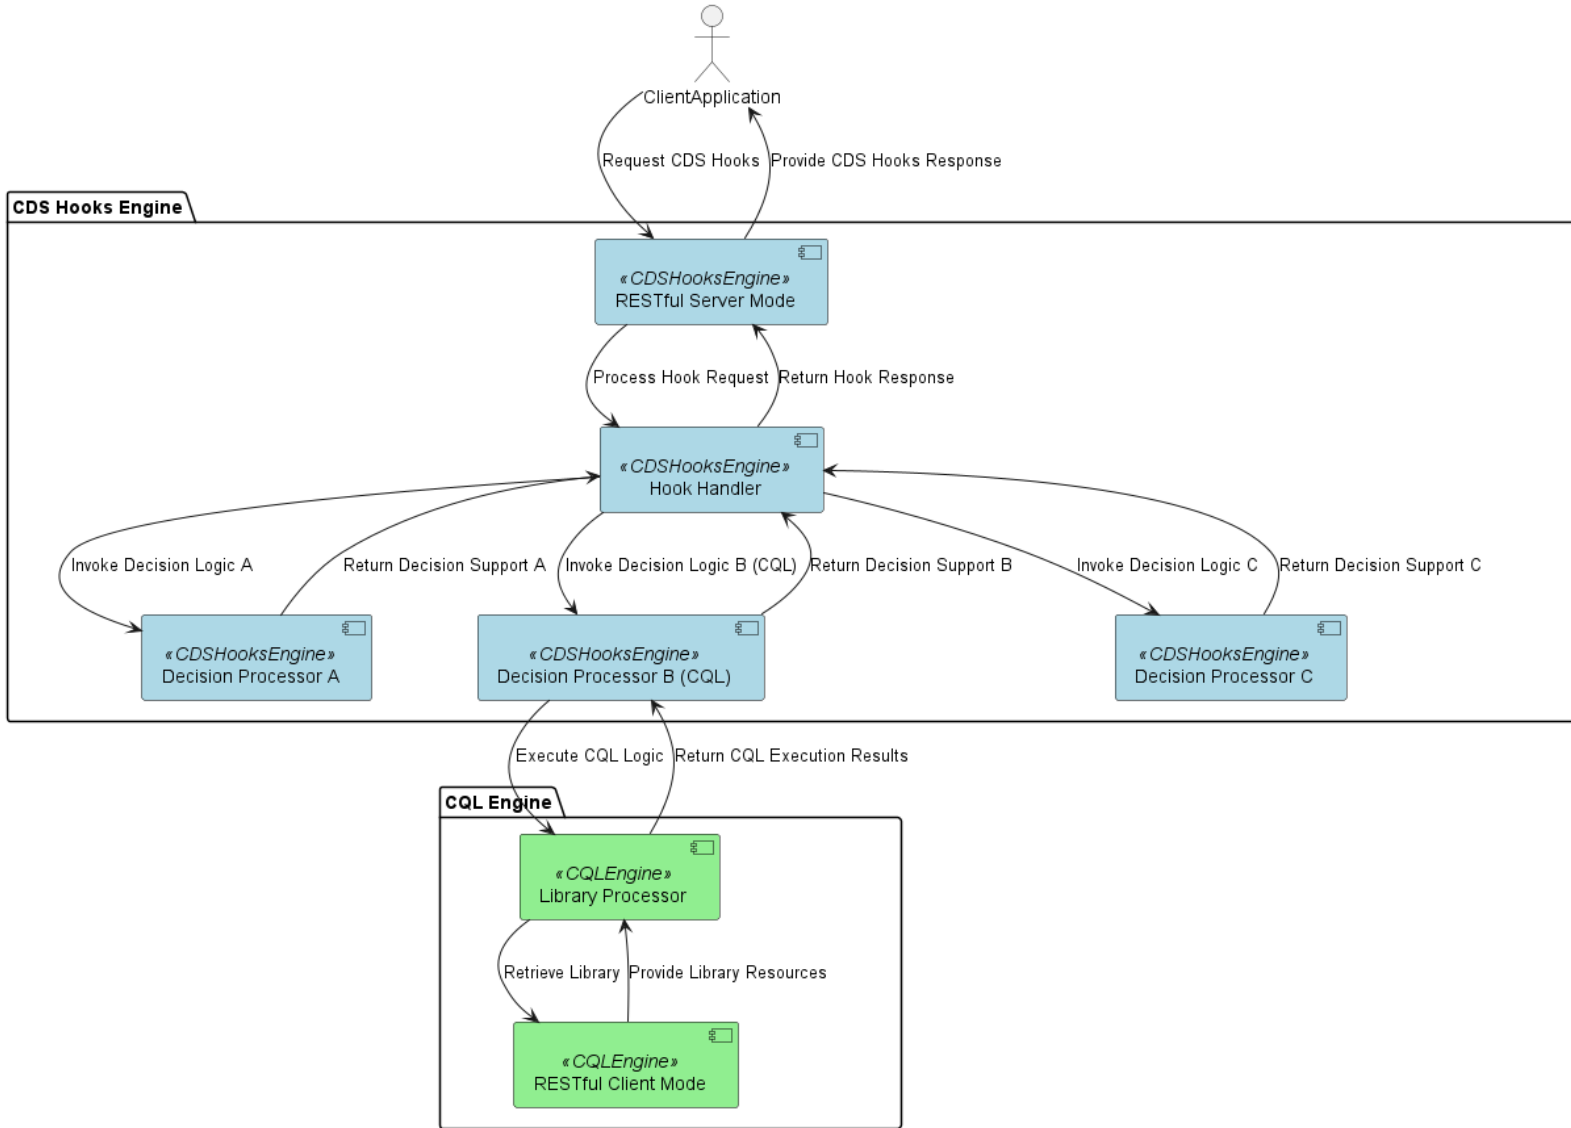


Figure 20: CDS Hooks engine workflow

This component was created in the context of the project as experimental development to test the display of clinical decision support that calculate the possible inclusion of the patient to display it in primary use softwares and in the future integrate AI assessment.

## 2.6 Future Recommendations

Looking ahead, further work could focus on consolidating and scaling the interoperability artefacts developed within FLUTE, including continued refinement of oncology data models, computable research protocols, and AI-related representation frameworks within HL7 FHIR. Strengthening alignment with initiatives such as HL7 Vulcan and CDISC will remain important to ensure regulatory compatibility and cross-domain interoperability. In addition, sustained collaboration through European frameworks such as PHOENIX and engagement with the evolving EHDS governance structures will be essential to promote adoption, implementation guidance, and long-term sustainability of FLUTE-derived oncology related specifications.

### 3 Conclusions

Through its contributions to the FLUTE project, HL7 Europe has strengthened the interoperability foundations of the project and positioned its outputs within a sustainable European and international standardisation trajectory. By aligning FLUTE artefacts with HL7 FHIR, engaging in oncology data modelling initiatives, supporting research artefact structuring, and contributing to imaging and AI-related discussions, HL7 Europe ensured that project results are reusable, standards-based, and compatible with emerging regulatory and cross-border frameworks.

These efforts not only supported FLUTE's immediate technical objectives but also enhanced its long-term impact by embedding its outcomes within recognised governance structures and active standardisation communities. As a result, FLUTE's contributions extend beyond the project lifecycle, reinforcing Europe's capacity for interoperable, federated, and secure health research.