



KODJIN
by Edenlab

Edenlab

The Semantic Foundation for AI-Ready Healthcare Data

Kseniia Nikolaienko — Head of Data & AI, Edenlab
HL7 Norway FHIR fagforum — 26 August 2026



“We have all this data in FHIR
now — can we just point an LLM
at it?”



FHIR solved exchange — not analytics

A FHIR resource is a self-contained unit of clinical meaning — exactly right for exchange. Point an analytical question, or an LLM, at millions of them, and the same richness becomes the obstacle:

Observation (FHIR JSON)

```
{
  "resourceType": "Observation",
  "code": {"coding": [{
    "system": "http://loinc.org",
    "code": "4548-4" }]},
  "subject": {"reference": "Patient/8291"},
  "component": [{
    "valueQuantity": {
      "value": 9.4, "unit": "%" }},
    ... references, versions, extensions
  ]
}
```

nested · reference-linked · one of millions

structural

layer

observations (analytics-ready)

patient_id	date	concept	value
8291	2026-03-14	HbA1c	9.4
4417	2026-03-12	HbA1c	7.1
1058	2026-03-11	HbA1c	10.2

flattened · terminology normalised · historised ·
streaming, seconds behind the FHIR server

one row per fact — queryable at population scale



Even ordinary questions have no obvious answer

Neither of these can be answered from raw FHIR without clinical modelling decisions — decisions that must be made deliberately, once, and not improvised inside every query:

“When was this patient exposed to this medication?”

MedicationRequest

what was prescribed

MedicationStatement

what the patient says they took

MedicationAdministration

what was actually given

exposure
= ?

three resources, three levels of truth — each with its own timing model

“When did this condition start?”

onsetDateTime

recordedDate

abatementDateTime



when it began —
if anyone knows

when someone
wrote it down

when it resolved —
if captured at all

each may be a date, an age, a period — or simply absent.

Which one anchors your cohort? A clinical modelling decision.

Between FHIR and the LLM: three engineered layers



01

A structural layer – a streaming transformation that turns nested FHIR resources into flattened, enriched, historised structures, available for analytics in seconds.

02

A semantic layer **of governed clinical and business concepts** — the thing that makes natural-language access safe, consistent, and explainable.


03

A privacy architecture that separates reasoning from data: the model works with metadata and has no access to your data, the query engine executes query — inside your governed infrastructure.



Making FHIR Analytics-Ready: The Structural Layer


Our pipeline rests on three ideas: stream every change as it happens, flatten it declaratively, and enrich it on the way in — so the data is always current



one change in the FHIR server → an analytics-ready, meaningful row — seconds later



Making FHIR Analytics-Ready: Architecture





The Semantic Layer: Where Meaning Lives

A governed, machine-readable catalogue of your organisation's concepts — patient, admission, episode of care, readmission — each formally defined once, instead of re-decided inside every report.


What this means in practice:

- 01 Concepts, not columns – “poorly controlled diabetics” resolves to the cohort definition your clinical team signed off — every time
- 02 Definitions carry their logic – tables and columns, code sets, relationships, business rules — everything a query needs travels with the concept
- 03 Tractable to build – your top ~100 concepts cover most real questions — governance work, not an ontology moonshot
- 04 Governed and versioned – clinical informaticians and analysts sign definitions off; every change is tracked, so yesterday's answer is reproducible today



From question to governed query

The semantic layer changes the LLM's job. It no longer writes queries against raw tables — it interprets questions against definitions your organisation controls:





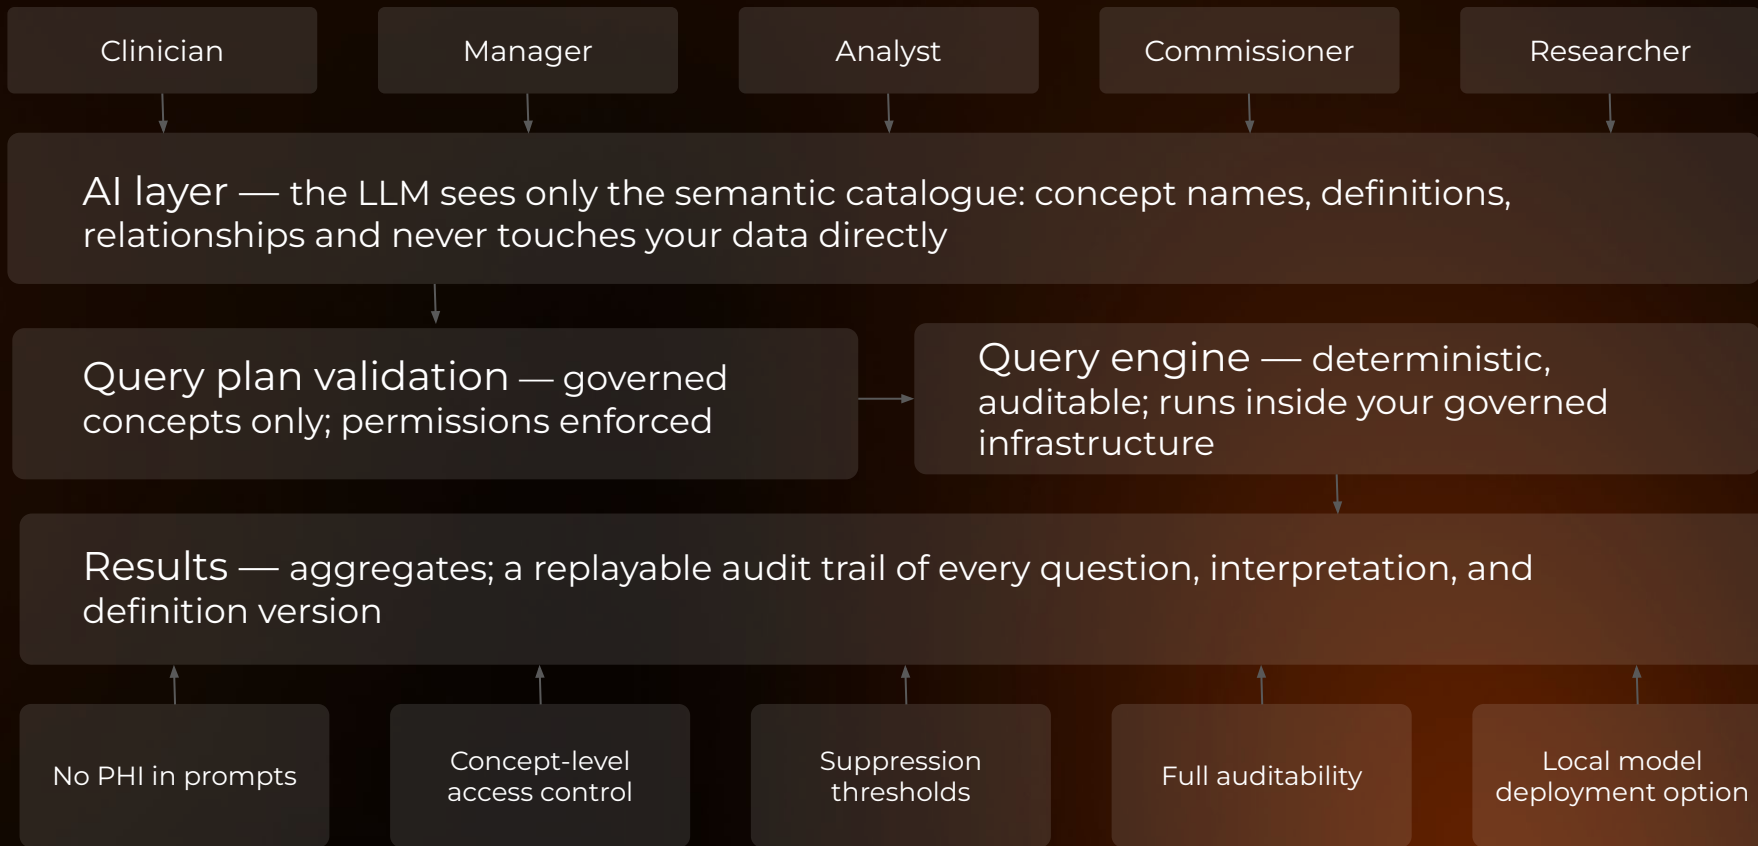
Why This Is a Safety Mechanism

The single biggest risk of natural-language analytics is silent misinterpretation. Governed concepts remove the improvisation:

Four reasons:

- 01** It constrains the model – questions map onto validated concepts; the query is generated deterministically from definitions. The model chooses from a menu — it doesn't cook
- 02** It makes answers consistent – same question, same number, however it is phrased — and every answer shows which concepts, filters, and time window were used
- 03** It encodes knowledge – arrhythmia and heart failure both count as cardiac events; gestational diabetes doesn't belong in this cohort — curated, not guessed from column names
- 04** It unlocks privacy – a model that understands your concepts does not need your data — descriptions are enough.

Privacy-by-Design: Separate Reasoning from Data





Demo:

conversational analytics
synthetic data, real
architecture



01

Exchange-ready is not AI-ready — between FHIR and any language model sits an engineered path: streamed, flattened, enriched, historised.

02

The semantic layer is the safety mechanism — governed concepts turn natural language from improvisation into validated interpretation.

03

Separate reasoning from data — the model sees metadata; the engine sees patient data inside your governed estate. Ask any vendor: what exactly does the model see?

04

Privacy-by-design is an architecture, not a policy document — design it in now, while these systems are being procured and built.

Thank you — questions?

Kseniia Nikolaienko

Head of Data & AI, Edenlab

 k.nikolaienko@edenlab.io

 kseniianikolaienko