




KELENO | HEALTHCARE INTEROPERABILITY CASE STUDY

FROM WEEKS TO MINUTES

Building an AI-native interface engine for healthcare interoperability

Plain English in. Production-grade HL7, FHIR, and X12 flows out — connecting EHR, lab, pharmacy, imaging, and payer systems in about 30 minutes, not 12 weeks.

Table of Contents

1. The Business Challenge.....	3
2. Why an AI-Native Interface Engine, Now.....	3
3. The Strategic Vision.....	3
4. The Solution.....	3
5. How It Works — Prompt to Production.....	4
6. Legacy Migration & Paper Workflows.....	4
7. Security by Design — Zero PHI in the Core.....	4
8. From Weeks to Minutes — The Impact.....	5
9. The Role of Keleno.....	5
10. Case Study at a Glance.....	5
Contact & Usage Policy.....	7

1. The Business Challenge

A US-based healthcare interoperability startup came to Keleno with a familiar, expensive problem: every EHR, lab, and payer system it needed to talk to depended on interface engines that were slow, costly, and run by a shrinking pool of HL7, FHIR, and X12 specialists.

A single interface routinely took six to twelve weeks to ship. Annual licensing ran into six figures before lock-in costs. And when something broke in production, almost nobody on the team could trace what data had gone where, or why.

The company's ambition wasn't to build another point-to-point integration. It was to remove interoperability as a bottleneck entirely — for itself first, and eventually for every health system and vendor it would need to connect to.

At a glance

- Industry: Healthcare Technology / Interoperability
- Technology: AI-native interface engine, agentic orchestration, HL7 v2, FHIR, X12
- Delivery: Keleno's India engineering team, working directly with the client's US-based product team
- Customer: US-based healthcare interoperability startup (anonymised for this case study)

2. Why an AI-Native Interface Engine, Now

In September 2026, OpenAI connected ChatGPT for Healthcare directly to Epic's EHR — read-only, permissioned, and gated by each health system's existing record-access rules. It was a clear industry signal: the next phase of healthcare AI isn't about smarter chat. It's about governed access to clinical data that already exists.

That governed access depends entirely on interfaces most patients — and most software buyers — never see. Every EHR, lab, and revenue-cycle system depends on interface engines to move data safely between systems, and those engines have historically been slow, expensive, and dependent on scarce specialist talent.

That was exactly the gap Keleno's client set out to close: building the layer that makes governed, AI-ready data exchange possible for any health system — not only the ones with the budget and time to wait months for a custom integration.

3. The Strategic Vision

The product vision was simple to state and hard to build: describe an integration once, in plain English, and have it built, validated, and shipped to production — without a specialist team and without a quarter on the calendar.

Conversational → Validated → Versioned → Production. In minutes, not months.

4. The Solution

Keleno's engineering team designed and built the interface engine around four principles: conversational authoring, self-healing validation, complete versioning, and a security architecture that keeps patient data out of the platform's own core.

The solution combined six layers:

1. A plain-English authoring layer that turns a written integration request into a build plan.
2. An orchestrated, multi-stage agentic workflow that generates, deploys, and validates the resulting flow.
3. A flow registry that versions every integration for rollback, diffing, and replay.
4. A vendor-pack library that encodes legacy EHR conventions for faster migration.
5. An AI-driven OCR pipeline that turns paper and fax-based inputs into structured HL7/FHIR.
6. A split-plane security architecture that keeps PHI out of the platform's control plane entirely.

5. How It Works — Prompt to Production

A user describes the integration they need in plain English — for example:

```
"Build an HL7 ADT interface from Epic to our data lake. Validate every message and map MRN to patient_id."
```

From there, an orchestrated, 11-node agentic workflow — five of the nodes LLM-powered — takes over: it listens for intent, generates the interface configuration, validates it against real schemas, and promotes it to production.

If validation fails, the workflow doesn't stop and wait for a human. The LLM reads the validation errors, corrects the offending nodes, and re-validates automatically, bounded by retry caps and deterministic guardrails so the system fails safe rather than looping indefinitely.

Every resulting integration is versioned in a flow registry: teams can roll back to any prior version, diff changes between versions, and replay transactions against historical flow definitions for audit or debugging.

6. Legacy Migration & Paper Workflows

Not every customer starts from a blank slate. The engine supports two paths for existing complexity:

- **Legacy flow import** — existing interfaces from engines such as Mirth Connect, Apache NiFi, or commercial platforms import directly, with routing intact, and can then be extended conversationally rather than rebuilt from scratch.
- **Vendor packs** — for teams migrating off legacy EHR platforms, pre-built packs encode the data conventions of Meditech, Allscripts, Cerner, NextGen, and eCW, so a migration starts from a known schema instead of a blank one — with full per-record reconciliation reporting.

A separate AI-driven OCR pipeline handles the paper side of healthcare that system-to-system interfaces can't reach: inbound faxes, scanned referrals, and PDF lab results are classified, field-extracted, validated against schema, and emitted as clean HL7 or FHIR directly into the target system.

7. Security by Design — Zero PHI in the Core

The architecture separates control from data. The control plane — identity, audit, and AI orchestration — never touches PHI; it only ever sees metadata and identity signals. The data plane,

where PHI actually lives and where flows execute, deploys directly into the customer's own cloud environment.

That split means protected health information never has to leave a customer's own infrastructure to be processed — a property of the architecture itself, not a policy layered on top of it after the fact.

8. From Weeks to Minutes — The Impact

~~12 weeks~~ → **30 minutes** prompt to production

Engineered outcomes the platform was built to deliver:

- **Speed** — an integration that took six to twelve weeks the legacy way now takes roughly thirty minutes, prompt to production.
- **Specialist dependency** — plain-English authoring and self-healing validation reduce reliance on scarce HL7/FHIR/X12 specialists for routine integration work.
- **Operational transparency** — versioned, diffable, replayable flows replace black-box interface behavior, so 2am incidents are traceable rather than mysterious.
- **Migration cost** — vendor packs let legacy EHR migrations start from a known schema instead of a blank slate, with reconciliation reporting built in.
- **Coverage** — the AI+OCR pipeline folds paper-based inputs into the same structured pipeline as system-to-system interfaces, closing a gap that governed EHR-native AI tools cannot reach.

These describe the engineered capability Keleno delivered. Customer-specific savings and volumes were not independently audited for this case study.

9. The Role of Keleno

Keleno's team owned design and engineering of the interface engine end to end — from the agentic orchestration architecture and the security model to the vendor-pack library and the OCR pipeline — working as an embedded extension of the client's own engineering team.

The engagement combined two of Keleno's core disciplines: AI & Agentic Engineering (LLM orchestration, self-healing agentic workflows) and IT Consulting & Systems Integration (HL7/FHIR/X12 domain modeling, legacy platform migration) — brought together under one delivery team rather than handed across separate vendors.

10. Case Study at a Glance

Customer	US-based healthcare interoperability startup
Industry	Healthcare Technology / Interoperability
Challenge	Slow, expensive, specialist-dependent interface engines
Solution	AI-native, conversational interface engine
Core standards	HL7 v2, FHIR, X12

Architecture	Split-plane design — control plane / data plane — zero PHI in the core
Orchestration	11-node agentic workflow, 5 LLM-powered, self-healing validation
Legacy support	Vendor packs — Meditech, Allscripts, Cerner, NextGen, eCW
Paper workflows	AI + OCR pipeline for faxes, scanned referrals, and PDF lab results
Speed	~30 minutes prompt-to-production vs. 6–12 weeks legacy
Keleno's role	End-to-end design & engineering, embedded delivery team

Contact & Usage Policy

Contact

For partnership and technology enquiries:

Website: <https://www.keleno.com>

Email: info@keleno.com

Phone: +91 88483 58534

Restricted Usage & Confidentiality

This case study is intended for Keleno website publication and approved stakeholder communication. It describes an engineered solution and platform capability; it does not constitute independently audited customer performance results. The customer is intentionally anonymised. No part may be reproduced, adapted, distributed, or used for commercial purposes without prior written permission from Keleno.

Any speed, cost, or effort figures referenced are descriptions of engineered platform capability and must be independently validated for a specific deployment, data volume, and operating environment.

© 2026 Keleno. All rights reserved.
