

ML LABS INTELLIGENCE

# AI ECG Platform for Clinical Operations

A medical device company needed real-time ECG analysis that could meet regulatory scrutiny and run inside existing clinical workflows. ML LABS built the production platform.

CASE STUDY

Author Omar Trejo

Date 2026-01-07

ML LABS

[mlabs.com/intel/ai-ecg-management-platform](https://mlabs.com/intel/ai-ecg-management-platform)

---

HeartSciences needed a production-grade cloud backend for MyoVista Insights, their vendor-neutral ECG management platform. The system had to process ECG signals in real time, support AI-ECG algorithm inference from multiple model providers, integrate bidirectionally with a major EHR vendor via HL7 and FHIR, and meet healthcare regulatory requirements on AWS — all while supporting multi-tenant clinical operations across US and UK deployments.

ML LABS designed and built the entire cloud-native backend, delivering the platform that now powers HeartSciences' end-to-end ECG ecosystem — from multi-vendor device ingestion through async AI inference to clinical workflow integration with signed reports and usage-based billing.

## The Problem

Building a medical device backend is not a normal software project. Every design decision carries regulatory weight, and the gap between "working prototype" and "production system regulators will accept" is where most teams stall. The compounding factor here was interoperability: the platform had to ingest ECGs from multiple device manufacturers, each with their own DICOM encoding conventions, and integrate with hospital EHR systems that expected standards-compliant HL7 messaging with specific segment configurations.

- Multi-vendor ECG format parsing across devices from multiple major manufacturers with measurement normalization differences
- Bidirectional HL7 messaging with the EHR vendor requiring order workflows and result delivery with diagnosis segments
- Regulatory constraints on data handling, auditability, and clinical report lifecycle
- Data protection compliance across the entire data pipeline with per-organization tenant isolation

- No existing backend infrastructure to build on

## Multi-Vendor ECG Ingestion

The platform had to be genuinely vendor-neutral — not just in marketing language, but at the DICOM parsing level. Devices from multiple major manufacturers each encode 12-lead ECG waveforms and measurements differently within their DICOM objects. Amplitude values arrive in different units, lead ordering varies, and measurement segments use vendor-specific structures.

ML LABS built a normalization layer that extracts waveform data and clinical measurements from each vendor's DICOM format into a unified internal representation. Deduplication at the storage layer prevents redundant waveform copies at ingestion — when a clinic uploads the same ECG through multiple pathways (common with file share ingestion from clinic networks), the system recognizes the duplicate and links rather than re-stores. This kept storage costs predictable as ingestion volume scaled across sites.

**The hardest interoperability problems are not protocol-level — they are the subtle measurement normalization differences between vendors that produce clinically incorrect values if handled generically.**

## AI Inference Architecture

The AI layer processes ECGs through multiple model providers — including AI model providers for rhythm analysis and cardiac function models for ejection fraction assessment. Each model has different latency characteristics, availability patterns, and billing implications.

ML LABS built an async inference architecture that decouples ECG submission from AI result delivery. The system handles provider unavailability gracefully, prevents duplicate billing when providers fail mid-processing, and supports per-organization model enablement — allowing HeartSciences to roll out new algorithms to specific sites without affecting others.

## EHR Integration Layer

EHR integration required bidirectional HL7 messaging — not just reading patient data, but participating in the clinical order workflow. The platform receives unsolicited order messages when a physician requests an ECG interpretation, processes the study through AI inference, and returns result messages with observation segments containing RR intervals, axis measurements, and clinical findings. Diagnosis segments carry billing codes that downstream hospital systems depend on for reimbursement.

On the FHIR side, the platform implements per-organization authentication that prevents cross-tenant data access — even if credentials are compromised, the system enforces organizational boundaries at every access point. This dual-protocol approach (HL7 for clinical workflow, FHIR for data access and launch context) reflects how EHR integrations actually work in production, not how they're described in documentation.

## Clinical Report Workflow

Clinicians don't just view AI results — they sign, confirm, and lock clinical reports through a lifecycle that carries legal and regulatory weight. ML LABS built a report workflow that enforces clinical accountability at every stage — once a report is signed, every subsequent modification is tracked with full audit trail visibility. The system ensures that no change to a signed clinical document goes unrecorded.

PDF generation handles the visual complexity of clinical ECG reports across multi-page layouts that maintain clinical readability. The system supports mobile viewing on iPad and iPhone with touch interactions for clinical measurement tools, allowing physicians to review and sign reports from wherever they are.

## Multi-Tenant Operations

The platform operates as a true multi-tenant system with per-organization configuration controlling everything from enabled AI models to HL7 field mappings to invoice pricing. New hospitals can be onboarded with site-specific configuration without touching the codebase.

The billing system implements usage-based invoicing that handles the edge cases making healthcare billing complex — partial months, mid-cycle onboarding, per-site minimums, annual caps, and tiered pricing. Multi-environment deployment across TST, PRD-US, and PRD-UK ensures that configuration changes are validated before reaching production clinical workflows.

## Security Posture

A platform handling PHI under regulatory oversight cannot treat security as a checklist exercise. ML LABS remediated 18+ security findings across the system, spanning credential management, data exposure prevention, injection protection across ingestion pathways, and authentication hardening. The security work was integrated into the delivery cadence rather than deferred to a separate hardening phase — each finding was addressed within the sprint where it was identified.

// Security in a medical platform is not a phase – it is a constraint that shapes every API surface, every data access pattern, and every integration endpoint from the first commit.

## Results

MyoVista Insights is now in production as HeartSciences' core platform — connecting ECG hardware from multiple manufacturers, AI-ECG algorithms from multiple providers, and clinical tools into one coherent ecosystem. The system processes ECGs in real time, maintains full regulatory audit trails, and operates reliably under production loads with automated scaling across US and UK deployments. Clinical reports generated by the platform achieved a 97% acceptance rate across participating institutions, meeting regulatory documentation standards without manual correction.

Clinicians use the platform to access ECGs, AI results, and prior studies in a single integrated view — including directly within the EHR. Health IT teams get standards-based interoperability without custom interfaces. A research SDK provides

programmatic access to ECG data across all ingested formats, enabling HeartSciences to support clinical research without building separate data export pipelines. A natural language search capability lets clinicians find studies using clinical language rather than navigating complex filter interfaces.

## First Steps

If your organization is building a regulated medical device or health-tech platform, the critical decision is whether to treat compliance as a design constraint or a post-hoc gate.

1. **Start with regulatory requirements as architecture inputs.** Build compliance into the data model, the deployment pipeline, and the monitoring layer from the first commit. The cost of retrofitting is always higher.
2. **Design for multi-vendor interoperability from day one.** Normalizing data formats after the fact introduces subtle clinical accuracy bugs that are expensive to find and dangerous to miss.
3. **Treat EHR integration as a first-class subsystem.** HL7 and FHIR integration patterns are well-documented in theory but full of per-organization edge cases in practice — budget for the reality, not the specification.

## Practical Solution Pattern

Build regulatory compliance and interoperability into the architecture from the start rather than treating them as integration layers bolted on later. Invest in vendor-neutral data normalization, async AI inference with billing safeguards, and

bidirectional EHR messaging with per-organization configuration. Instrument every clinical workflow transition for auditability, and design the multi-tenant model to support site-level capability rollout without code changes.

This works because the cost structure of regulated systems is fundamentally front-loaded — architectural decisions made in the first months determine whether the system can scale to new sites, new AI models, and new EHR partners without rework. Organizations that ship these platforms fastest concentrate authority in a single technical decision-maker who holds the full regulatory, clinical, and engineering picture simultaneously, rather than distributing it across committees that optimize for coordination over delivery. If you need to decide whether a regulated workflow is ready for build, a **Strategic Scoping Session** can pressure-test scope, constraints, and next steps before engineering work starts.

## References

1. U.S. Food and Drug Administration. **Software as a Medical Device (SaMD)**. *Regulatory Reference*, 2024.
2. HeartSciences. **MyoVista Insights**. *HeartSciences*, 2024.
3. HL7 International. **FHIR Standard**. *HL7*, 2024.
4. HL7 International. **HL7 v2.x Messaging Standard**. *HL7*, 2024.
5. NIST. **Healthcare Security Rule Guidance**. *National Institute of Standards and Technology*, 2024.



**ML LABS**

Custom AI Systems for High-Value Workflows

[mllabs.com](http://mllabs.com)