

# The Regulatory Guardrail Build for Healthtech

A 14-day architectural build to eliminate dangerous regulatory drift by anchoring your AI coding agents in a deterministic understanding of your SaMD and clinical regulatory domain.

## THE EXECUTIVE DILEMMA

- **The Regulatory Compliance Trap:** Large Language Models excel at generating code quickly, but they operate on probabilities. They do not know your unwritten clinical invariants, ePHI isolation boundaries, or strict regulatory constraints. When unleashed on a clinical, SaMD, or mission-critical healthtech codebase with only natural language requirements, development agents confidently guess. They invent parallel clinical workflows, state transitions or unauthorized data paths, creating rogue compliance logic. This taxes your most senior engineers with untangling hallucinated clinical logic and ePHI boundaries during code review.
- **The Compliance Illusion:** Transitioning to an AI-assisted SDLC increases velocity, but FDA, MDR, and EU AI Act auditors will never accept a probabilistic LLM prompt as a system safety guarantee. If engineering teams scale these tools without deterministic boundaries, the resulting regulatory drift introduces unacceptable patient-safety risks and existential legal liability.
- **The Solution (Deterministic Compliance Verification):** I extract the compliance logic currently relying on manual code reviews or probabilistic LLM validation and codify it into a formal regulatory domain model and executable verification gates. Plugged into your agentic development workflow, it triggers agentic verification checks with mathematical determinism, forcing development agents to self-correct against HIPAA and MDR reality before a single line of code is written.

## ENTERPRISE OUTCOMES

- **Audit-Ready AI Compliance:** By translating clinical invariants, ePHI access states, and HIPAA/MDR rules into an executable domain boundary, you generate deterministic proof that your AI cannot violate system safety. You can safely scale agentic coding workflows while generating a formal Compliance Bill of Materials (CBOM) for agentic execution, providing the exact deterministic proof auditors require.
- **Intercepting Regulatory Tech Debt:** By establishing the formal regulation-grounded domain model as your executable single source of truth, conflicting specifications, unauthorized data handling, or hallucinated state transitions trigger immediate failures and feedback during agentic development. This blocks the agent from drafting a non-compliant pull request.
- **Protected Engineering Bandwidth:** By codifying regulatory knowledge into an executable and formal baseline, your developers and autonomous agents are aligned with the same regulatory understanding. Instead of senior engineers burning hours acting as manual AI safety-checkers, the drop-in agentic hooks enforce your MDR/HIPAA boundaries upfront, protecting your team's cognitive load.
- **Immediate Talent ROI:** Stop forcing your most expensive asset to act as manual AI compliance-checkers. If a €150k+/year senior engineer spends just 5 hours a week untangling hallucinated compliance logic, they are wasting over €18,750 of capacity annually. By intercepting these errors deterministically upstream, this fixed-scope build eliminates the PR bottleneck, paying for itself entirely within weeks through reclaimed developer velocity.

## CORE DELIVERABLES

*This is a capability transfer, not a consulting retainer. At the end of the 14-day build, your team owns a production-ready compliance guardrail and the methodology to scale it internally without ongoing vendor dependency.*

- **The Formal Domain Model & Invariant Baseline:** A formal, machine-readable domain model defining a single, high-value regulatory slice (a bounded context) and associated regulatory invariants. This domain-specific model strictly maps the exact clinical rules, ePHI state transitions, and regulatory constraints your agents must respect to safely navigate that specific area of your architecture.
- **The Drop-In Agentic Guardrail:** I don't install heavy proprietary software. I deliver the exact plugin components (agentic skills, hooks, and tools) that force your existing agentic setup to comply with the domain model and force agentic self-correction. Operating quietly in the background, this guardrail acts as a local pre-flight check, grounding agentic development actions in your regulatory reality deterministically. Crucially, it automatically produces a formal Compliance Bill of Materials (CBOM) for every approved execution, establishing a deterministic audit trail of why the architecture is compliant.
- **The Architectural Evolution Runbook:** You are not trading AI babysitting for heavy maintenance. To ensure your team retains 100% independent authority, I deliver a definitive maintenance runbook for your senior engineers and compliance officers. It provides the step-by-step workflow for your team to easily update and extend the domain boundaries as your clinical business rules evolve.

## DELIVERY MECHANICS

- **Targeted Extraction (90-Min Session):** A single, high-efficiency working session with your leads to pinpoint high-leverage targets and existing ground-truth rules (e.g., FHIR contracts, clinical data schemas, ePHI boundaries).
- **Modeling & Guardrail Engineering:** I independently extract and formalize your clinical rules into a formal domain model and regulatory invariants, tailoring the automated verification hooks to your existing agentic development environment.
- **Air-Gapped Delivery:** I do not require access to your live CI/CD pipelines, internal networks, or sensitive source code. The compiled domain model and guardrail hooks are delivered as self-contained artifacts, bypassing InfoSec bottlenecks.
- **Interactive Handover:** We run self-contained validation test fixtures inside your local environment, demonstrating the guardrail actively intercepting hallucinated regulatory logic and proving determinism across real agent workflows.
- **Fixed-Scope Investment:** €19,500 (Excl. VAT). Structured as an Operational Velocity Investment, sized specifically to fall within standard VP of Engineering / CTO discretionary spending limits for immediate activation.
- **Next Step:** Book a 30-minute Alignment Call to map this capability to your existing SDLC bottlenecks.

## ABOUT TONI RUOKOLAINEN, PH.D.

I am the Founder and Principal Architect specializing in Domain-Specific Modelling, Formal Methods, and AI Enablement. I hold a Ph.D. in Computer Science with a focus on model-driven engineering, complex system modeling and formal methods, backed by over 15 years of deep technology research. I help highly regulated healthtech and SaMD engineering organizations accelerate their product delivery with AI-assisted SDLC built on an uncompromising, deterministic regulatory grounding.