

Ironwell Health

Your guardian for trusted AI in healthcare

Healthcare organizations face pressure from every side: rising costs, staffing shortages, new regulations, and a flood of AI tools promising to fix it all. You don't need another vendor. You need a partner who stands with you, protects what matters most, and helps you move forward with confidence. That's who we are.

What we do:

We build, evaluate, and safeguard AI and interoperability platforms in healthcare technology.

Who we serve:

Healthcare provider, payers, vendors, along with venture capital and private equity firms invested in healthcare companies.

AI Solutioning and Build

From first idea to lasting impact, with you at every step.

Discovery. Not sure where to start? You're not alone, and you don't have to figure it out by yourself. We'll walk alongside your team to find the business processes that are ready for agentic optimization, so your first investment goes where it matters most.

Readiness Assessment. Before we build anything, we make sure you're set up to succeed. We'll evaluate your data, systems, security posture, and team readiness, and uncover the roadblocks that stall AI projects before they cost you time and money.

ROI and Business Case. We'll help you make the case with confidence. For each opportunity, we'll set baseline metrics and project the return, giving your leadership clear numbers to say yes to.

Pilot. We'll build a pilot for one of up to three prioritized use cases from Discovery. It's a safe, contained way to test assumptions and prove real impact on your business processes before you commit to a full rollout.

Compliance by Design. Your patients' trust is sacred, and we treat it that way. Every solution we build includes HIPAA and PHI safeguards, Business Associate Agreement review for model and cloud vendors, and human oversight wherever clinical or financial decisions are made. Protection isn't added at the end. It's built in from day one.

Go Live. When your pilot proves its value, we'll deploy your agentic solution to production and train your team to manage it with confidence. You'll own it, and you'll understand it.

Managed Operations. Our commitment doesn't end at launch. We'll keep watch over your AI solutions, monitoring performance, accuracy, cost, and reliability, and tuning them each month so they keep delivering. You'll get regular reports and peace of mind.

Scale and Expand. Success should build on itself. We'll map a roadmap for expanding to new use cases on the same foundation, so every agent after the first arrives faster and costs less.

AI Development Workflows

Empowering your engineers, safeguarding your codebase.

AI Coding Policy and Governance. Every great team needs clear guardrails. We'll help you set an acceptable-use policy that defines which tools are approved, what code and data can be shared, and how intellectual property and licensing are protected, so your developers can innovate without risk.

Tooling. We'll partner with your software engineers to bring AI into your development process in a way that scales and actually supports your developers instead of frustrating them. Your team's expertise stays at the center. AI just helps them do more of what they do best.

Secure AI-Assisted Development. Speed should never come at the cost of security. We'll put protections in place, including secret scanning, dependency and license checks, and controls that keep PHI and credentials out of AI prompts, so your organization stays protected as your team moves faster.

CI/CD Guardrails. We'll help you add guardrails and automated testing for AI-written code. You'll know you're not checking in code that hides technical debt waiting to explode later. Your codebase stays clean, stable, and built to last.

Code Review Enablement. We'll train your reviewers to spot the common weaknesses in AI-generated code, and set up AI-assisted review tools that strengthen human judgment rather than replace it.

Productivity Measurement. You deserve proof, not promises. We'll set baselines and build dashboards that track cycle time, defect rates, and review load, so you can see whether AI is truly improving your team's work.

Process Documentation and Training. We'll help you document your updated SDLC and train your developers to run your AI-enhanced workflow on their own. We're building your independence, not dependence on us.

AI Governance, Auditing and Compliance

Your protection against risk, so you can innovate without fear.

AI Inventory. You can't protect what you can't see. We'll find and catalog every AI tool and model in your organization, including the "shadow AI" your teams may have adopted on their own.

AI Risk Assessment and Audit. We'll review your AI systems against leading frameworks and regulations, including the NIST AI Risk Management Framework, ISO 42001, ONC HTI-1 transparency requirements, and emerging state AI laws, so you know exactly where you stand and what to strengthen. We can also test against your existing AI governance program.

Bias and Performance Testing. Every patient deserves fair, accurate care. We'll test your models across patient populations to make sure your AI serves everyone equitably.

AI Governance Program. We'll help you build a lasting governance program, with a committee structure, clear policies, an approval workflow for new AI tools, and board-ready reporting.

AI Vendor Due Diligence. Before you sign, we'll look closely on your behalf. We'll evaluate third-party AI products for data handling, transparency, and contract terms, so you know you're making the right decision.

RAG and Knowledge Platforms

Putting your organization's knowledge to work, securely.

Knowledge Assistant Build. We'll build secure, permission-aware AI assistants that give your teams instant, accurate answers from your policies, clinical guidelines, payer contracts, and procedures. Everyone sees only what they're authorized to see.

Content Readiness. Great AI starts with great information. We'll clean, structure, and organize your source documents, the step that most often decides whether a knowledge platform succeeds.

Evaluation and Grounding. Your teams need answers they can trust. We'll test for accuracy, citation quality, and hallucinations before launch and on an ongoing basis.

AI-Assisted Interoperability

Connecting providers, payers, and vendors, with integrity at every link.

Interoperability Assessment. We'll map your data flows and find gaps against FHIR, HL7, TEFCA, and CMS interoperability and prior authorization requirements, giving you a clear path to compliance and connection.

AI-Assisted Data Mapping. We'll use AI to speed up the work of mapping legacy formats, HL7v2, and FHIR, with expert human validation at every step so nothing gets lost in translation.

Payer-Provider Data Exchange. We'll build and modernize the workflows that connect your organization to its partners, including prior authorization, claims, and clinical data exchange, so information moves quickly, securely, and accurately.

AI Security

Standing guard over your AI systems.

Agent and LLM Security Testing. We'll test your AI systems the way an attacker would, probing for prompt injection, data leakage, and excessive agent permissions, so weaknesses are found by us and not by someone else.

Guardrail Implementation. We'll put strong protections around your production AI, including input and output filtering, PHI redaction, and access controls, keeping your patients' data and your reputation safe.

Fractional Leadership (vCTO and vCISO)

Executive expertise at your side, without the executive overhead.

Virtual CTO. Strategic technology leadership on your terms: roadmaps, architecture guidance, vendor selection, and AI strategy aligned with your business goals.

Virtual CISO. A dedicated security leader in your corner. We'll guide your security program, conduct HIPAA risk assessments, prepare you for HITRUST and SOC 2, and keep your board informed and confident.

Training and Enablement

Empowering your people to lead with confidence.

Executive AI Briefings. Clear, practical guidance that helps your leadership and board understand AI's opportunities and risks, and make decisions with confidence.

Staff AI Literacy. Role-based training for clinical, revenue cycle, and administrative teams, so everyone can use AI safely, effectively, and with confidence.

vCTO Services

Strategy and roadmap

- Technology strategy tied to business goals, with a 12–36 month roadmap
- AI strategy: where AI fits, build vs. buy, use-case prioritization
- Architecture reviews and modernization planning (cloud, EHR integration, data platforms)

Vendor and spend management

- Vendor selection, RFP support, and contract review for EHR, RCM, AI tools, and cloud
- Managing the technology budget and cutting overlapping or unused tools

Engineering and delivery oversight

- Overseeing internal teams or outsourced developers
- Setting development standards, MLOps/LLMOps practices, and release processes
- Technical due diligence for investors, acquisitions, or partnerships

Data and interoperability

- Data strategy and governance: who owns what data and how it moves
- Interoperability planning: FHIR/HL7, payer-provider data exchange, TEFCA, and CMS interoperability rules

Leadership

- Board and executive reporting on technology

- Hiring support for technical staff, and eventually helping hire a full-time CTO
- vCISO Services

Governance, risk, and compliance

- HIPAA Security Rule risk assessments, which are required and frequently cited in OCR enforcement
- Writing and maintaining security policies and procedures
- Readiness for HITRUST, SOC 2, ISO 27001, and NIST CSF
- Compliance mapping for AI rules: the ONC HTI-1 algorithm transparency requirements, state AI laws, NIST AI RMF, and ISO 42001

Security program management

- Security roadmap and maturity scoring
- Third-party and vendor risk management, including reviewing Business Associate Agreements (BAAs) and AI vendors' data handling
- Overseeing identity and access management

Incident readiness

- Incident response plans and tabletop exercises, with emphasis on ransomware, which is a major threat in healthcare
- Breach response coordination, including the HIPAA breach notification process
- Business continuity and disaster recovery planning

Testing and monitoring

- Coordinating penetration tests and vulnerability management
- AI-specific security testing: prompt injection, data leakage, and model abuse in agentic systems
- Security awareness training and phishing simulations

Executive and board support

- Board-level reporting on cyber risk
- Support with cyber insurance applications and renewals

- InfoSec questionnaire / RFP support

Our Technical Expertise:

AWS, Bedrock, Agentcore, Langchain, LangGraph, LangSmith, Threads, Python, Node.js, SQL, Postgresql, NoSQL, DataBricks, Bitbucket, Github, Generative AI, Agentic AI, Large Language Models (LLMs), ChatGPT, Claude, LangChain, Amazon Bedrock, Amazon Bedrock AgentCore, n8n, Salesforce, Salesforce Agentforce, Salesforce AppExchange, Apex, SaaS, Python, TensorFlow, Machine Learning, Supervised Learning, Random Forest, AI-Driven Data Scoring, Prompt Engineering, Human-in-the-Loop AI, Hallucination Management, Agentic Drift Monitoring, SIEM, Proactive Alerting, AWS, Amazon EC2, Amazon RDS, Amazon VPC, AWS Fargate, AWS CloudFormation, Microservices, Cloud-Native Architecture, Microsoft Azure, On-Premises Infrastructure, Cloud Migration, Network Architecture, iPaaS, Mirth Connect, SSIS, Node.js, PostgreSQL, Microsoft SQL Server, MySQL, Stored Procedures, Database Indexing and Tuning, VB.NET, RESTful APIs, API Integration, FHIR, HL7, C-CDA, XML, JSON, X12, Middleware Integration Development, Results Interfaces, Demographics/Charges/Scheduling Interfaces, Epic, Meditech, Cerner, McKesson, McKesson Patient Folder, Allscripts, HealthQuest, Iatric, NIST, HIPAA, HITRUST, SOC 2 Type II, Security Controls, Security Policy and Procedure Design, Security Hardening, Security-First Design, HIPAA Incident Response, Disaster Recovery, Business Continuity, Acceptable Use Policy, Infrastructure Diagramming, AI Governance, LLM Governance, AI Data Governance, Data Platform Architecture, Product Architecture, Provider Directory Data Management, Practitioner Data Management, Credentialing and Recredentialing, Technical Due Diligence, M&A, Post-Merger Integration, Exit Readiness, Value Creation, SaaS Transformation, Board Reporting, C-Suite and PE Stakeholder Engagement, Process Documentation, Training Development