

WHITE PAPER

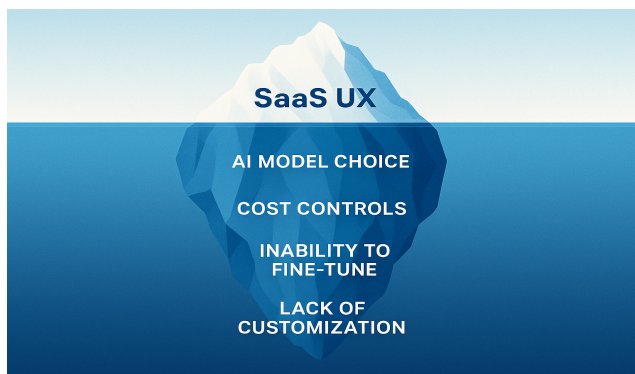


From Opaque to Open: Re-Engineering SaaS AI

How MSOs, Health Plans, and TPAs Regain Data Control & Model Freedom

WHY READ THIS (2 MIN)

- AI is now woven into just about every SaaS platform you license—yet you rarely control which model runs, how it's trained, or what it costs.
- CEOs feel trapped by **rising per-seat “AI uplift” fees** with no visibility into declining wholesale model costs.
- COOs struggle to **fine-tune models on proprietary workflows** or swap in best-of-breed models as needs evolve.



THE BLACK-BOX PROBLEM

Embedded AI Reality	Strategic Impact
Single, vendor-chosen model hidden behind an API	No leverage when accuracy or cost is sub-optimal
Terms claim broad rights over prompts & outputs	Risk of IP leakage and regulatory scrutiny
Fine-tuning limited to generic fields (if at all)	Can't encode unique processes, compliance rules
No roadmap for model updates or portability	Future innovation gated by the SaaS vendor



Data Ownership & Customization Barriers

WHO REALLY OWNS YOUR DATA & IDEAS?

- **Prompt & Output Capture:** Many SaaS vendors reserve the right to retain prompts and AI outputs for “service improvement.”
- **Shadow Training Sets:** Those retained records may end up in the vendor’s own foundation-model fine-tuning—outside your control.
- **Audit & HIPAA Risk:** For healthcare entities, co-mingling PHI or proprietary payer rules in a vendor’s opaque pipeline threatens compliance.

WHY YOU CAN’T CUSTOMIZE (YET)

1. **Static Model Choices** – Most vendors lock you to their partnership (e.g., GPT-4 Turbo today, Anthropic, Gemini, unknown tomorrow).
2. **Limited Vector Stores** – You cannot link your domain data without exporting it to the vendor’s closed index.
3. **Governance Gaps** – Explainability (XAI) tools seldom exposed, leaving risk teams guessing.

THE BLACK-BOX PROBLEM

Scenario	Pre-Intervention	Post-Intervention (open AI layer)
MSO automating provider onboarding	18 min manual compliance review per file	4 min with RAG-powered checklist
Regional Health Plan call-center QA	9% escalations flagged late	2% escalations

(Figures reflect conservative composites of recent Serious Development engagements.)



A Flexible, Plug-and-Play Architecture

THE ESCAPE PLAN: BUILD AN OPEN AI CONTROL LAYER

1. **Agentic Orchestration**
Micro-agents select the best model per task (summarization, reasoning, code) in real time.
2. **Model Control Panel**
Swap GPT-4o for Claude Haiku or fine-tuned Phi-3 without touching downstream workflows.
3. **Private Vector DB & RAG**
Keep sensitive provider, member, or claims data inside your VPC; enrich prompts securely.
4. **Explainability (XAI) Hooks**
Feed SHAP/LIME scores to risk teams; store citations for audit.
5. **Deep-Research Mode**
Long-context analysis over policy, clinical, or contractual docs—no vendor seat surcharge.

MINI CASE STUDY – LARGE REGIONAL PROVIDER GROUP

- **Problem**
Embedded AI in HRIS did not allow for a changing criteria set in order to match ideal candidates. The ability to fine tune and control prompt/criteria is not available in their HRIS.
- **Solution**
SD designed a customized plug-and-play AI solution that works with our client's existing HRIS, incorporates other important data, and allows for ongoing fine tuning of prompts and criteria.
- **Results (2 mo)**
Candidate match accuracy increased by 42%, manual review time decreased by 45%, and HR team productivity improved significantly, with ongoing incremental gains from weekly fine-tuning. Better yet, they are improving the quality of their hires.



ROI Checklist, Glossary & Next Steps

RAPID ROI CHECKLIST

Tick off where an open AI layer can deliver value in **< 12 months:**

✓ Checkpoint	Typical 12-mo Gain	Notes
Replace per-seat SaaS AI fees with usage-based model billing	25-50 % AI cost reduction	AI wholesale prices falling ~35 % YoY
Fine-tune on proprietary workflows (onboarding, sentiment analysis, etc)	40-200% increase in productivity	Drives time savings
Enable model swap-outs without re-implementing workflows	Future-proof innovation / better results	Avoid 6-figure re-licensing events

GLOSSARY (SELECT TERMS)

- **Agentic Approach:** Network of specialized AI agents collaborating on sub-tasks.
- **MCP (Model Control Panel):** an open standard to enable AI assistants to connect seamlessly with external data sources and tools. This protocol facilitates more relevant and context-aware responses from AI models by standardizing how applications provide context to large language models (LLMs)
- **RAG (Retrieval-Augmented Generation):** Method that injects domain documents into model context for grounded answers.
- **Vector Database:** Stores embeddings for fast similarity search over large corpora.
- **XAI (Explainable AI):** Techniques that make model outputs understandable to humans and regulators.

Where to Go from Here

Opaque SaaS AI doesn't have to dictate your costs—or your roadmap.

Organizations that **layer a flexible, transparent AI foundation on top of existing systems** unlock faster innovation, avoid vendor locking, tighter governance, and measurable savings.

Connect with a custom software partner who understands healthcare workflows and AI—**like Serious Development**—who can tailor, implement, and evolve that layer around your existing applications, data, workflows, and compliance requirements.