

REVSPRING

AI Posture Statement

How RevSpring uses AI on behalf of our clients,

and what we will and will not do with it.

FOR CLIENTS, AUDITORS, AND PARTNERS

Version 1.0 | Effective Date Pending

A note to our clients

AI is now a normal part of how patient billing and engagement work gets done. RevSpring uses it. Our peers use it. The companies you compare us against use it. The question that matters is not whether AI is involved — it is whether the company involving it can tell you, clearly and without hedging, what it does, who it acts for, and how it is held accountable when it gets something wrong.

This document is our answer.

We have designed our AI program around a single principle: every AI agent we deploy at RevSpring is treated as a digital twin of a human role. It has a name. It has a sponsor. It operates within a published scope. Every action it takes is logged against both its own identity and the identity of the human it acts on behalf of. If you ask us who took an action, we can show you both.

The rest of this document spells out what that means in practice — the commitments we make, the boundaries we hold, and the evidence we will provide when you ask for it. We have written it in plain language because you should not need a translator to understand how your data is being handled.

If anything here raises questions or does not match what your contract requires, please raise it with your RevSpring contact. We would rather have the conversation than discover the gap later.

— *RevSpring Leadership*

How we think about AI

RevSpring operates AI under a model we call the Digital Twin Operating Model. The premise is straightforward.

The premise

Every AI agent inside RevSpring is paired with a human sponsor. The agent — the "twin" — can only do what the sponsor could do, and only within an explicit, narrower scope that the sponsor and their manager have approved in advance. The twin has its own identity in our systems. Its actions are logged against both its identity and the sponsor's. When something goes wrong, accountability is unambiguous because two named parties stand behind every action.

Why we built it this way

Most AI policies in our industry treat AI as a tool. A tool has a user, but no accountability of its own — when a tool produces a bad outcome, the conversation defaults to a fight about whether the user was at fault or the vendor was at fault. That is a poor foundation for regulated services.

Treating AI as a digital twin moves the conversation. The twin is accountable. The sponsor is accountable. RevSpring is accountable. The audit log shows all three. This is the model that fits the language regulators, auditors, and clients already use to evaluate workforce access — because it is, in practical terms, a workforce model.

What this means for your data

When a RevSpring twin touches data that belongs to your organization or your patients, it does so under the same boundary as the human sponsor. Your existing BAA, MSA, and data processing agreements with RevSpring already define that boundary. The Digital Twin model does not expand it. It tightens it, by adding a second layer of identity, scope, and audit on top of the controls you already have.

Eight commitments

These are the commitments RevSpring makes to every client. They apply to every AI agent we deploy, regardless of vendor, model, or use case.

COMMITMENT 01

We will never sell or share your data to train external AI models.

Client data, patient data, and any identifiable information processed on your behalf is never used to train models maintained by RevSpring or by any third party. Our contracts with model providers prohibit retention of your data for training purposes.

COMMITMENT 02

Cardholder data is out of scope for every twin.

No AI agent at RevSpring is permitted to receive PAN, CVV, or full track cardholder data. This boundary is enforced at the data layer, not the prompt layer — meaning a twin cannot bypass it even if instructed to. Twins see masked indicators (last-4 only) and never the full payment instrument.

COMMITMENT 03

PHI is handled under the same BAA boundary as our human workforce.

Twins that touch PHI operate through a PHI-cleared model gateway with no-retention contractual terms. The sponsor of any PHI-touching twin holds active HIPAA training. The BAA you signed with RevSpring covers our twins as workforce extensions of their sponsors.

COMMITMENT 04

Your data does not leak into another client's twin.

For client-facing twins, we operate per-tenant instances or enforce strict context isolation. We red-team this boundary on a recurring basis. If we ever discover a tenant cross-contamination event, you will be notified per the breach terms in your contract, regardless of whether regulatory notification thresholds were crossed.

COMMITMENT 05

Every twin action is logged against two identities.

The agent that took the action, and the human sponsor it acted on behalf of. We retain these logs per the retention terms in your contract and produce them on request for audits, incident investigations, or consumer rights requests under applicable privacy laws.

COMMITMENT 06

Twins draft. Humans send.

For any output that reaches you or your patients — outreach messages, account notes, statements, eligibility narratives — a human sponsor reviews and submits the work. We will not move to autonomous outbound communication on your account without an explicit, documented agreement from your team.

COMMITMENT 07

We disclose AI involvement where the law or your contract requires it.

Where applicable law, your contract, or your specific policy requires that an AI's involvement be disclosed to a patient or consumer, we disclose it. We do not present AI-generated communication as having been authored by a specific named human.

COMMITMENT 08

We retire twins that drift.

If a twin's measured performance declines, if its scope is no longer appropriate, or if its sponsor leaves, we follow a documented remediation path: reduce scope, retrain or replace the model, reassign the sponsor, or retire the twin. Orphaned AI identities are not allowed to persist.

What we will not do

We think it is as important to be specific about what we will not do as it is to be specific about what we will. The boundaries below apply across our AI program, regardless of client request.

- We will not deploy AI agents that operate autonomously on your account without an explicit, written, tier-appropriate authorization from both RevSpring and your organization.
- We will not process cardholder data through any AI system, regardless of model provider or contractual claim.
- We will not use your data — or your patients' data — to train, fine-tune, or improve any AI model that operates outside the boundary of your services with us.
- We will not allow an AI agent to take a final action on a payment, refund, write-off, or account closure without a human in the chain.
- We will not impersonate a specific named RevSpring employee, your employee, or a clinician in AI-generated communication.
- We will not run AI evaluations, A/B tests, or pilots on your patient population without your prior written approval.
- We will not bypass any breach notification, audit cooperation, or evidence production obligation in your contract because the action in question was taken by an AI rather than a human.

Regulatory alignment

RevSpring's AI program is designed to align with the regulatory frameworks our clients operate under. Brief summaries below; full control mappings are available on request under NDA.

Framework	RevSpring AI Posture
HIPAA	AI agents touching PHI are workforce extensions of BAA-covered sponsors. PHI-cleared model gateway with no-retention contract terms. HIPAA training required for any sponsor of a PHI-touching twin.
PCI DSS	Cardholder data is outside the scope of every twin. The boundary is enforced at the data layer, not the prompt layer. Twins see masked indicators only.
SOC 2	Twin actions are part of our Common Criteria control set. Dual-identity audit logs, scope manifests, and the AI Registry serve as evidence artifacts. Available to clients under SOC 2 reliance letters.
TCPA & State Comms Laws	AI drafts; humans send. Consent and channel rules continue to be enforced through our existing communications platform, not through the twin layer.
State Privacy (CCPA, etc.)	Consumer rights requests include twin-generated artifacts. Our AI Registry supports the lookup. Automated decisioning disclosures are made where applicable.
FDCPA / Reg F	Where collections workflows fall under FDCPA, AI-drafted communications are reviewed by trained human collectors before any external send. Required disclosures are not generated by AI.

If your jurisdiction or contract introduces additional obligations not listed above, we will document the specific control mapping in a client-specific addendum to this Posture Statement.

What you can ask us for

The following are available to clients on request. Request channels and turnaround times are defined in your services agreement.

Audit log extracts

Dual-identity logs of every AI action taken on your account or your patients' accounts, scoped to a date range and use case you specify. Provided in machine-readable format suitable for ingestion into your own audit systems.

Scope manifests

The published scope of any twin operating on your account, including data sources read, actions taken, and approval boundaries. Useful for your own due diligence reviews.

AI Registry entries

The full Registry record for any twin operating on your account: name, sponsor, manager, tier, charter date, last review date, and current status.

Incident reports

Documentation of any AI-related incident affecting your account, including root cause, remediation, and changes made to prevent recurrence. Provided per the incident notification terms in your contract.

Annual posture review

An annual review session with your account team and our AI Platform leadership covering year-over-year changes, new use cases, and any policy updates. We recommend pairing this with your existing SOC 2 review cadence.

Client-specific addenda

Where your specific regulatory environment, internal policy, or risk posture requires it, we can produce a client-specific addendum to this Posture Statement that documents the additional commitments we make to your organization.

In closing

AI is going to be part of how this industry operates. The companies that handle it well will be the ones who can answer hard questions with specific evidence — not the ones with the most polished marketing about innovation.

RevSpring's bet is that you would rather work with a company that can show you a named twin, a named sponsor, and a logged action than one that can show you a logo and a press release. The Digital Twin Operating Model is how we make that bet operational.

We expect this Posture Statement to evolve. Regulatory expectations will sharpen. Client expectations will sharpen. Our own program will sharpen. We will publish updates as they happen, and we will tell you what changed and why.

Thank you for reading this far. The fact that you did is the reason we wrote it this way.

— End of Posture Statement v1.0 —