SpencerFane®

A Case Study in Privacy Risks: The Character.AI Bankruptcy

Recently, I wrote several articles on the Health Care Governance Gap.[1] In short, the central problem is straightforward: Many health care organizations have readily adopted artificial intelligence (AI), but fewer have clearly assigned responsibility for vetting it prior to deployment, validating its outputs, monitoring its performance, documenting oversight, and establishing procedures for intervening when it fails. Legal may interpret the rules, compliance may monitor, information technology (IT) may implement, and cybersecurity may protect the infrastructure – but clinicians may ultimately bear the consequences. Yet, in many settings, no one has fully operationalized the controls that link these functions. That disconnect is the "Governance Gap," and it matters more in health care than in most sectors because a system's output can directly affect patient lives, reimbursement, and trust. It is already costing organization billions.

A further emerging risk is the loss of patient trust when AI tools seem more authoritative than they really are. A patient-facing tool that explains symptoms, provides educational material, or assists with navigation can add real value. A tool that appears to present itself as a licensed clinician, or that leaves patients unable to tell automated guidance apart from professional care, is a different matter altogether. That is why Pennsylvania initially sued Character.AI: The company's chatbots have posed as physicians and given medical advice, even fabricating medical license numbers.

Character Technologies Inc., the company behind the companion chatbot platform Character.AI, recently became the first consumer-facing AI company to file for Chapter 11 bankruptcy. The filing exposed something most governance frameworks were never built to handle. Chat logs from more than 20 million users, including

more than 2 million of them from minors, showed many treated the chatbots as therapists and confidants – discussing their mental health, family struggles, relationships, and most private secrets. Privacy advocates have called it one of the most troubling data scenarios in recent memory – not because the data was stolen, as I typically see in my own practice – but because it is now being treated as a financial asset inside a bankruptcy proceeding that was never designed to account for conversational records or the AI models trained on them.

A proposed licensing deal with a major entertainment company reportedly carves out the chat logs themselves, but legal observers note that it doesn't resolve much: The data is likely already embedded in the underlying model, thereby making it difficult, if not impossible, to cleanly separate it out the way a company might sell or delete an ordinary customer list. Courts have floated the use of privacy ombudsmen, typically used to anonymize names, emails, and birthdates, as a safeguard. But that mechanism wasn't built for intimate confessions from its most vulnerable users: children. Commentators question whether it can meaningfully protect people who never anticipated their most personal disclosures becoming a line item in a liquidation.

This isn't just a cautionary tale about one company. This is a preview of a governance failure health care organizations should recognize immediately, because it mirrors a problem already unfolding in clinical settings: AI adoption is outrunning the frameworks meant to hold it accountable.

The Data Problem Doesn't End When the Vendor Does

What the Character.AI bankruptcy adds to this picture is a dimension many health care organizations haven't fully confronted or appreciated: what happens to sensitive data, and the models trained on it, when the company holding it no longer controls its own future. Clinical AI vendors, scribes, and patient engagement platforms collected information that can be even more sensitive than a chatbot conversation – and that information is protected by HIPAA and other legal obligations that don't simply evaporate because a vendor goes bankrupt, gets acquired, or shuts down.

But as this case study illustrates, bankruptcy law treats data as an asset to be valued, licensed, or sold; not necessarily as a liability bound by the original promises made to the people who generated it. If no health care-specific contingency exists in a vendor contract, an organization may find itself with little leverage over what happens to patient data once that vendor's business model fails. Waiting for a court-appointed ombudsman to sort it out after the fact is not a strategy: It's an admission that governance should have happened earlier.

Building Audit Trails Before You Need Them

If an AI scribe fills in documentation, a utilization model influences approvals or denials, or a language model drafts patient communications, organizations need an audit trail strong enough to answer basic questions later: What did the system do? What data did it rely on? Who reviewed it? Was the output accepted, modified, or rejected? Without that record, governance becomes performative – and when a vendor relationship ends badly, whether through bankruptcy, breach, or simple failure, organizations without that record may be unable to defend their process to regulators, payors, courts, or patients.

The regulatory backdrop is only getting more specific. The U.S. Food and Drug Administration revised its Clinical Decision Support Software Guidance in January 2026, clarifying when provider-facing software functions fall outside device regulation and when oversight still applies, particularly where clinicians can't independently review the basis for a recommendation. The European Union AI Act continues pushing health care AI toward a high-risk governance model centered on risk management, data governance, human oversight, logging, and post-market monitoring, with penalties reaching €35 million or 7% of global turnover. Regulators are signaling, clearly and repeatedly, that health care AI will be judged not just by what it can do, but by how responsibly it's governed – including after a vendor relationship ends.

What Health Care Organizations Should Do Now

Health care organizations should treat AI as enterprise risk, not an isolated innovation initiative: Inventory which tools are in use, classify them by risk, and document ownership across legal, compliance, IT, security, privacy, and clinical

leadership. Require documented human review wherever AI output could materially affect a patient, a treatment decision, or a claim. Build real audit trails capturing inputs, outputs, reviewers, and escalation decisions. And, critically, negotiate data-stewardship and contingency terms into every AI vendor contract upfront – addressing deletion, transfer, and retraining obligations if that vendor is ever acquired, dissolved, or placed in bankruptcy.

Character.AI's bankruptcy is a warning shot to others about what happens when sensitive data becomes an asset with no clear rules attached. Health care organizations have an opportunity – and thus an obligation – to make sure they never find themselves relying on a bankruptcy court to protect what their patients trusted them with in the first place.

Key Takeaways

- AI in health care creates real value but also real operational, regulatory, and patient-safety risk.
- The biggest failure point is often not the model itself, but the absence of clear accountability for review and oversight.
- Patient-facing AI must never blur the line between automated assistance and licensed care.
- Vendor bankruptcy can turn sensitive patient data into a stranded, high-risk asset – plan for it contractually before it happens.
- Regulators are increasingly focused on traceability and governance, not just innovation.

This blog post was drafted by Christine Chasse, an attorney in the Plano and Dallas, Texas, offices of Spencer Fane. For more information, visit www.spencerfane.com.

———————

[1] See also: AI in Health Care: The Governance Gap and In Healthcare, an AI Mistake Can Cost a License or a Life

*Click **here** to subscribe to Spencer Fane communications to ensure you receive timely updates like this directly in your inbox.*