

AI IN CHILD WELFARE:

WHY THE FUTURE DEPENDS ON
GETTING THE FOUNDATION RIGHT



How to Verify a Vendor Before You Trust Them With Your Clients' Data

These aren't bureaucratic checkboxes. Each one represents a real risk that has materialized at real agencies. Work through them in order.

Look for compliance seals — then verify they're real

HIPAA, CJIS, and SOC 2 badges should appear in the footer of any vendor's website. But a badge is easy to fake. Click through to the documentation. **A real HIPAA claim means a signed BAA is available. Real CJIS compliance means documented alignment with the FBI's security policy. Real SOC 2 means an audit was conducted by an accredited third party.**

Find the Trust Center

Reputable platforms serving sensitive industries publish a dedicated Trust Center — a section of their site that documents their security posture, certifications, data handling practices, and compliance status. **If a vendor doesn't have one, ask why. If they can't answer clearly, you have your answer.**

Read the Terms and Conditions — specifically the data use clauses

Most people skip the T&Cs. Don't. The two sections that matter most:

- 1. Whether the vendor uses your data to train AI models.**
- 2. How long they retain your data after you leave.**

These are the clauses that determine what happens to your clients' information. Read them before you sign.

Ask for the audit reports

Don't ask if they're compliant. Ask for the documentation. **Request the signed BAA template, the CJIS compliance statement, and the most recent SOC 2 Type II report.** A vendor with nothing to hide will send these without hesitation. A vendor that hedges, deflects, or offers a one-page summary instead of the actual report is telling you something.

Ask about model training explicitly

Say it plainly: "Will any content our staff submits — including recordings, transcripts, or case-related text — be used to train or improve your AI models?" **Get the answer in writing.** For platforms handling child welfare data, the only acceptable answer is an unambiguous no.

What AI Can Actually Do in This Field

AI isn't magic, and it isn't a replacement for the humans doing this work. But there are places where it's already adding real value.

- **Documentation and case notes**

Social workers spend a disproportionate share of their time on documentation — time that comes directly out of face-to-face contact with families. AI tools that help structure, draft, or summarize case notes can meaningfully reduce that burden. Less time with a keyboard. More time with people.

- **Forensic interview skill development**

AI-powered avatar tools — simulated children that interviewers practice with — lead to measurable skill improvements: more open-ended questions, better technique, fewer protocol deviations. The improvement curve is faster than traditional training alone, and it doesn't require access to a supervisor for every session.

- **Post-interview analysis**

AI tools can analyze transcripts of forensic interviews to flag question types, identify moments where the interviewer may have introduced suggestion, and surface patterns across multiple sessions. A supervisor who previously had to watch hours of footage can now get a data-informed first pass in minutes.

- **Mitigating bias in investigations**

When asked to generate alternative hypotheses for a given scenario, AI systems have outperformed expert investigators, naive raters, and psychologists on the number of alternatives generated. In a field where confirmation bias can have devastating consequences, a thinking partner that doesn't share your assumptions is worth taking seriously.

- **Court preparation**

Forensic interviewers, social workers, and CPS investigators are using AI to simulate cross-examination, anticipate defense challenges, and stress-test their recollection of cases before taking the stand. Done correctly — with no identifying client information in the tool — this is a low-risk, high-value use.

Where It Gets Complicated

- **Predictive tools and racial bias**

Several high-profile studies have documented that predictive risk assessment tools in child welfare encode and amplify racial and

socioeconomic disparities. A tool trained on historical data will learn the patterns of historical decision-making – including the biases embedded in it. “The algorithm flagged it” is not a defensible substitute for human judgment.

- **The documentation trap**

AI-drafted case notes are only as good as the information that goes into them. A tool that produces fluent, professional-sounding documentation from thin inputs can create a false sense of completeness. Courts don’t just need notes that read well – they need notes that are accurate, specific, and grounded in direct observation.

- **Dependency and deskilling**

AI can accelerate the feedback loop in forensic interviewing – it shouldn’t replace it. The goal is interviewers who are better at their jobs. Not interviewers who can’t practice without an AI in the room.

