CASCADE AI CONSULTING

Using AI Safely in Social Services

A Practical Guide to HIPAA, Client Privacy & De-Identification for AI Use

A Free Resource from Cascade AI Consulting
cascadeaiconsulting.com | cole@cascadeaiconsulting.com

This guide is written specifically for homeless services providers, behavioral health organizations, community action agencies, and other social services nonprofits that want to use AI tools to reduce administrative burden while protecting client privacy and maintaining regulatory compliance.

Part 1: Why This Guide Exists

AI tools like ChatGPT and Claude can dramatically reduce the time your staff spend on documentation, grant reporting, and administrative work. Organizations using AI-assisted documentation report saving 15-20 hours of staff time per week. But in social services, we work with some of the most vulnerable people in our communities, and their information demands the highest level of protection.

The good news: **you can use AI safely and effectively without compromising client privacy.** The key is understanding what the rules actually require, where the real risks are, and how to build simple habits that protect your clients while letting your staff benefit from these tools.

This guide covers the regulatory landscape (HIPAA, HMIS, 42 CFR Part 2), explains the real-world risks in plain language, provides a step-by-step de-identification protocol your staff can follow, and gives you organizational policies you can adopt today.

Who This Guide Is For

- Executive Directors and leadership deciding whether to adopt AI tools
- Program Directors managing staff who handle client documentation
- Compliance officers and HIPAA Privacy Officers evaluating AI risk
- Frontline staff who want to save time on documentation but are worried about privacy
- IT and operations staff evaluating AI tool vendors

Part 2: The Regulatory Landscape

Before diving into practical steps, it helps to understand the regulations that govern how your organization handles client data. This section explains each regulation in plain language and what it means for AI use specifically.

2.1 HIPAA (Health Insurance Portability and Accountability Act)

HIPAA is the primary federal law governing the privacy and security of health information. It applies to covered entities (health plans, healthcare providers, and clearinghouses) and their business associates.

Does HIPAA apply to my organization?

If your organization provides any healthcare services, bills insurance, or receives healthcare funding, you are likely a covered entity under HIPAA. Many behavioral health providers and homeless services organizations that provide health-related services fall under HIPAA. Even if you are not a covered entity, following HIPAA-level privacy protections is best practice for any organization handling sensitive client data.

What HIPAA protects: Protected Health Information (PHI)

PHI is any individually identifiable health information held or transmitted by a covered entity. This includes information that relates to a person's past, present, or future health condition, the provision of healthcare, or payment for healthcare, when combined with identifiers that could identify the individual.

The 18 HIPAA Identifiers:

#	Identifier	Examples
1	Names	First name, last name, initials
2	Geographic data	Street address, city, zip code (first 3 digits OK if population >20K)
3	Dates	Birth date, admission date, discharge date, date of death, ages >89
4	Phone numbers	Any phone number
5	Fax numbers	Any fax number
6	Email addresses	Personal or work email
7	Social Security numbers	Full or partial SSN
8	Medical record numbers	MRN, chart numbers

#	Identifier	Examples
9	Health plan beneficiary numbers	Insurance ID numbers
10	Account numbers	Financial account numbers
11	Certificate/license numbers	Driver's license, professional licenses
12	Vehicle identifiers	License plate, VIN numbers
13	Device identifiers	Medical device serial numbers
14	Web URLs	Personal websites, social media profiles
15	IP addresses	Computer or network addresses
16	Biometric identifiers	Fingerprints, voiceprints, retinal scans
17	Full-face photographs	Photos or comparable images
18	Any other unique identifier	Case numbers, client IDs, any number that could identify someone

What HIPAA means for AI use:

Consumer AI tools like ChatGPT and Claude (free or Pro versions) are **not** HIPAA-compliant. They are not covered by a Business Associate Agreement (BAA), and data entered into them may be used for model training or stored on servers you do not control. This means you cannot enter PHI into these tools. However, you **can** use them safely by de-identifying information before entry (see Part 4).

> Some AI vendors do offer HIPAA-compliant, BAA-covered enterprise versions of their tools. If your organization uses one of these tools with a signed BAA, the data handling rules are different. Always verify BAA status with your vendor before assuming compliance.

2.2 HIPAA Security Rule Updates (2025-2026)

In January 2025, HHS proposed the first major update to the HIPAA Security Rule in 20 years. The proposed rule, expected to be finalized by mid-2026, has significant implications for organizations using AI:

- **AI systems that process ePHI must be included in risk analysis** and risk management activities

- **All security safeguards become mandatory** (eliminating the current distinction between "required" and "addressable" measures)
- **Organizations must maintain asset inventories** of all technology, including AI tools, that create, receive, maintain, or transmit ePHI
- **Annual compliance audits** will be required
- **Encryption** of ePHI at rest and in transit will be mandatory, not optional

Even though these rules are not yet final, organizations should begin preparing now. The direction is clear: HHS expects covered entities to treat AI tools with the same rigor as any other system that handles protected health information.

2.3 HMIS Data Standards

The Homeless Management Information System (HMIS) is governed by HUD data standards that impose specific privacy and security requirements on all Continuum of Care (CoC) funded organizations. HMIS data is considered personally identifiable information (PII) and must be protected accordingly.

Key HMIS privacy requirements for AI use:

- HMIS data must not be shared with unauthorized systems or third parties
- Client consent is required for data sharing beyond direct service provision
- Organizations must maintain data security plans that cover all systems handling HMIS data
- Any breach of HMIS data must be reported to the CoC and HUD
- AI tools that are not part of the approved HMIS ecosystem are considered unauthorized systems

The bottom line: HMIS data (client names, dates of service, housing status, income, demographic information) should never be entered into consumer AI tools in identifiable form.

2.4 42 CFR Part 2 (Substance Use Disorder Records)

If your organization provides substance use disorder (SUD) treatment services, 42 CFR Part 2 applies. This federal regulation provides even stronger privacy protections than HIPAA for SUD treatment records. Under Part 2, records cannot be disclosed without specific written patient consent, and re-disclosure is generally prohibited. Records protected under Part 2 should **never** be entered into any AI tool, even in de-identified form, unless you are certain the de-identification is complete enough to prevent any possibility of re-identification.

2.5 State Privacy Laws

Many states have enacted or are considering AI-specific legislation. As of 2026, over 250 AI-related bills have been introduced across more than 34 states. Key trends include requirements to disclose AI use, annual impact assessments, anti-bias controls, and record-keeping requirements. Oregon and Washington organizations should monitor state-level AI legislation in addition to federal requirements.

Part 3: Understanding the Real Risks

Fear about AI and privacy often comes from not understanding **where the actual risks are**. This section separates real risks from overblown fears so your organization can make informed decisions.

3.1 How Consumer AI Tools Handle Your Data

When you type information into a consumer AI tool (like the free version of ChatGPT or Claude.ai), here is what typically happens:

- **Your input is sent to the AI company's servers** over the internet (encrypted in transit)
- **The AI processes your input** and generates a response
- **Your conversation may be stored** on the company's servers for a period of time
- **Your data may be used to improve future AI models** (this varies by provider and plan level)
- **AI company employees may review conversations** for safety, quality, or abuse prevention

This is why entering identifiable client information is problematic: you are effectively disclosing that information to a third party (the AI vendor) without a BAA, without client consent, and without control over how the data is stored, used, or accessed.

3.2 The Real Risks vs. Overblown Fears

Concern	Real Risk Level	Why
Someone reads my AI conversation and identifies a client	LOW (if de-identified)	If you remove identifiers, there is nothing to link back to a real person
AI "remembers" my client across multiple conversations	LOW	Consumer AI tools generally do not retain memory between conversations (unless you enable that feature)
A data breach at the AI company exposes client information	MODERATE (if PHI entered)	This is a real risk if identifiable data is entered; mitigated by de-identification
AI-generated text contains incorrect information	HIGH	"Hallucination" is common; AI may add details not in your notes. Always review output.
Staff use AI without any guidance or policy	HIGH	Without clear guidelines, staff may inadvertently share PHI. Policy and training are essential.

Concern	Real Risk Level	Why
Regulatory action for using AI	LOW-MODERATE	HHS has not prohibited AI use; the risk is using it improperly with identifiable data

The most common and preventable risk is not a sophisticated data breach. It is a staff member copying and pasting a client's full name and date of birth into ChatGPT because nobody told them not to. This guide exists to prevent that.

3.3 Consumer AI Tools vs. Enterprise/HIPAA-Compliant AI Tools

Feature	Consumer AI Tools	Enterprise / BAA-Covered AI
Examples	ChatGPT (free/Plus), Claude.ai (free/Pro)	ChatGPT Enterprise/Team with BAA, AWS Bedrock, Azure OpenAI
HIPAA-compliant?	No	Yes (with signed BAA)
BAA available?	No	Yes
Data used for training?	Potentially yes	Generally no (opt-out)
Can I enter PHI?	NO	Yes (with safeguards)
Cost	Free or $20/month	$25-60/user/month+
Best for	De-identified documentation, general drafting, admin tasks	Clinical workflows with PHI, EHR integration, data analysis

Most social services organizations will start with consumer AI tools using de-identification, which is effective, free or low-cost, and requires no vendor contracts. As your organization matures in AI use, you may choose to invest in enterprise tools with BAA coverage for workflows that require PHI.

Part 4: The De-Identification Protocol

De-identification is the process of removing or altering information that could be used to identify a specific individual. When done properly, de-identified data is no longer considered PHI under HIPAA and can be shared with AI tools without violating privacy regulations.

The Golden Rule

Never put identifying information into consumer AI tools.

You de-identify → AI writes → You add identifiers back

4.1 What to Remove Before Using AI

Always remove the following before pasting into any consumer AI tool:

Information Type	Example	Replace With
Full names	"John Smith"	"client" or "individual"
Date of birth	"03/15/1985"	"late 30s" or "30s"
Street addresses	"1234 SE Hawthorne Blvd"	"SE Portland" or "east side"
Phone numbers	"503-555-1234"	[remove entirely]
Social Security Numbers	"123-45-6789"	[remove entirely]
Medical record numbers	"MRN 4567890"	[remove entirely]
Email addresses	"jsmith@email.com"	[remove entirely]
Specific dates of service	"December 24, 2025"	"recent contact" or [remove]
Unique identifiers	Case numbers, ID numbers	[remove entirely]
Photos or images	Any image of a client	[never upload to AI]
Names of family members	"His wife Sarah"	"client's partner"
Specific employers	"Works at Fred Meyer on 82nd"	"employed at a retail store"

4.2 What Is Generally Safe to Include

Information Type	Example	Why It's OK
Age range	"40s," "mid-30s," "elderly"	Not specific enough to identify
Gender	"male," "female," "non-binary"	Not identifying on its own
General location	"downtown," "east side"	Broad geographic area
Housing status	"unsheltered," "in shelter"	Status category, not location
General health status	"reported chronic pain"	De-identified health info
Services provided	"provided naloxone"	Your actions, not client identity
Client statements (paraphrased)	"expressed frustration with..."	Paraphrased, not verbatim quotes
Your observations	"appeared tired," "engaged"	Your clinical judgment
Plan / next steps	"scheduled follow-up"	Future actions, not identifying

4.3 The Stranger Test

Before pasting anything into an AI tool, ask yourself one question:

> ***"If a stranger read this, could they identify who I'm talking about?"***
>
> *If YES → Remove more identifying details*
> *If NO → You're ready to use AI*

The Stranger Test is simple, fast, and effective. It accounts for combinations of details that might not be identifying on their own but could identify someone when combined. For example, "a 37-year-old Native American woman with twins living in a shelter on the east side" might be identifiable even without a name, depending on the community size.

4.4 Before and After: A Real-World Example

BEFORE de-identification (DO NOT paste this into AI):

> "Met with John Smith, DOB 3/15/85, at his apartment at 1234 SE Hawthorne. He said he's been using heroin again since losing his job at Target. Gave him my card (503-555-1234) and scheduled follow-up for Thursday December 19th. Wife Sarah was present and seemed concerned."

AFTER de-identification (safe to paste into AI):

> "Met with client, late 30s male, at his apartment in SE Portland. He reported returning to heroin use following recent job loss. Provided contact information and scheduled follow-up. Client's partner was present and expressed concern about the situation."

Notice what changed: the name, date of birth, specific address, phone number, employer name, partner's name, and specific date were all removed or generalized. The clinical content is fully preserved.

Part 5: The 5-Step Safe AI Workflow

Step 1: Capture

During or right after a client contact, jot down key bullet points: what happened, what you did, what the client said (key points), and what's next. This can be bullet points on paper, phone notes, or a voice memo.

Step 2: De-Identify

Before opening your AI tool, scan your notes and remove or replace all identifying information. Use the table in Section 4.1 as your guide. Apply the Stranger Test.

Step 3: Generate

Open your AI tool (Claude.ai or ChatGPT), paste the appropriate prompt template, add your de-identified bullets, and generate the draft.

Step 4: Review and Edit

Always read the entire AI output. Check for accuracy (did AI add anything not in your notes?), appropriate tone, and any errors. AI "hallucination" is common; delete anything that was not in your original notes.

Step 5: Re-Identify and Submit

Copy the edited draft into your EHR or documentation system. Add back the client's name, date of birth, specific dates, and case numbers. Do a final review. Submit. Optionally, delete the AI conversation after you are done.

Pre-Use Checklist

Before every AI use:

- ■ Names removed or replaced with "client"
- ■ DOB removed or replaced with age range
- ■ Addresses removed or replaced with general area
- ■ Phone numbers removed
- ■ ID numbers removed (SSN, MRN, case numbers)
- ■ Passed the Stranger Test

After every AI use:

- ■ Read the entire AI output
- ■ Checked for accuracy (no hallucinated details)

- Edited as needed
- Added identifiers back in EHR/documentation system
- (Recommended) Deleted AI conversation

Part 6: Organizational AI Privacy Policy Template

This section provides a ready-to-adopt organizational policy for AI use. Replace bracketed sections with your organization's specific information.

6.1 Policy Statement

[Organization Name] is committed to protecting the privacy and confidentiality of all client information. As we adopt AI tools to improve operational efficiency, we will ensure that all AI use complies with HIPAA, HMIS data standards, 42 CFR Part 2 (where applicable), and all relevant state and federal regulations.

6.2 Approved AI Tools

Tool	Approved For	NOT Approved For	BAA Status
[e.g., Claude.ai]	De-identified documentation drafting, admin tasks	Any use with PHI or identifiable data	No BAA
[e.g., ChatGPT]	De-identified documentation drafting, admin tasks	Any use with PHI or identifiable data	No BAA
[e.g., Enterprise tool]	Clinical documentation with PHI (if BAA signed)	Sharing outside approved use cases	BAA in place

6.3 Required Staff Training

All staff who use AI tools must complete training that covers:

- The 18 HIPAA identifiers and what constitutes PHI
- The de-identification protocol and Stranger Test
- The 5-step safe AI workflow
- How to handle mistakes (report, don't hide)
- Approved vs. prohibited uses of AI tools
- How AI hallucination works and why review is mandatory

Training must be completed before staff use AI tools and refreshed annually or when policies change.

6.4 Prohibited Uses

Staff must NEVER use consumer AI tools to:

- Enter client names, dates of birth, Social Security numbers, or any of the 18 HIPAA identifiers
- Upload client photos, voice recordings, or documents containing PHI
- Enter HMIS data in identifiable form
- Enter substance use disorder treatment records (42 CFR Part 2)
- Make autonomous decisions about client eligibility, services, or safety
- Use AI tools on personal accounts for work-related client documentation
- Share AI-generated content containing client information via email or messaging

6.5 Incident Response

If a staff member accidentally enters identifiable client information into a consumer AI tool:

1. **Stop** immediately. Do not enter more information.

2. **Delete** the AI conversation. (In Claude: click three dots, select Delete. In ChatGPT: click three dots, select Delete chat.)

3. **Report** the incident to your supervisor and [HIPAA Privacy Officer / Compliance Lead].

4. **Document** what information was shared, when, and which tool was used.

5. **Assess** the risk level. The Privacy Officer will determine if this constitutes a reportable breach.

Create psychological safety around reporting. The goal is learning and improvement, not punishment. Staff who hide mistakes create greater organizational risk than staff who report them.

Part 7: Special Considerations for Social Services

7.1 Harm Reduction and Outreach Services

Harm reduction and street outreach programs often work with individuals who may be undocumented, involved in the criminal justice system, or engaged in survival behaviors. The stakes of accidental disclosure are especially high in these contexts. Extra care should be taken to remove:

- Specific locations of encampments or squats
- Details about illegal activity (drug use, sex work) that could be linked to an individual
- Immigration status or documentation status
- Specific physical descriptions that could identify someone in a small community
- Nicknames or street names that might be identifying

7.2 Behavioral Health Documentation

Mental health and substance use disorder records carry additional protections. When using AI for behavioral health documentation:

- Never enter diagnosis codes or specific diagnosis names combined with any identifying information
- Remove all references to specific medications (replace with "prescribed medication" or "medication management")
- Be extra cautious with psychotherapy notes, which have heightened protections under HIPAA
- 42 CFR Part 2 records require an additional layer of protection beyond standard HIPAA requirements
- When in doubt, consult your compliance officer before using AI with behavioral health documentation

7.3 Small Community Considerations

In smaller communities or specialized programs, de-identification requires extra care. Even without a name, a combination of details might identify someone. For example, in a program serving 30 families, "a 45-year-old Native American woman with four children currently in shelter" might be identifiable to anyone familiar with the program. In these cases:

- Use broader age ranges ("40s" rather than "45-year-old")
- Omit race/ethnicity unless clinically relevant to the note
- Use "client" without gender if the population is small enough for gender to be identifying

- Omit or generalize the number of family members
- Consider whether the combination of safe details creates an identifiable profile

7.4 Client Consent and Transparency

While de-identified data is not covered by HIPAA consent requirements (because it is no longer PHI), best practice in social services is to be transparent with clients about how your organization uses technology. Consider:

- Adding a statement to your intake materials: "Our organization may use AI tools to assist with administrative tasks. No identifying information is shared with these tools."
- Being prepared to answer client questions about AI use honestly and simply
- Including AI use in your organization's Notice of Privacy Practices if applicable
- Respecting client preferences: if a client objects to AI use in their documentation, accommodate that request

Part 8: Evaluating AI Vendors for Privacy & Compliance

As your organization considers AI tools beyond consumer versions, use this checklist to evaluate vendors:

Privacy & Compliance Questions to Ask Every AI Vendor:

- Will you sign a HIPAA Business Associate Agreement (BAA)?
- Is our data used to train your AI models? Can we opt out?
- Where is our data stored? (Geographic location and cloud provider)
- What is your data retention policy? Can we request deletion?
- Do you have SOC 2 Type II certification?
- What encryption standards do you use (at rest and in transit)?
- Who at your company can access our data? Under what circumstances?
- What is your incident response and breach notification process?
- Can we get a copy of your most recent security audit?
- Do you have a Data Processing Agreement (DPA) available?

Vendor Evaluation Scorecard:

Criterion	Score (1-5)	Notes
BAA available and signed	___	
Data not used for model training	___	
SOC 2 Type II certified	___	
Data stored in US / compliant jurisdiction	___	
Clear data retention and deletion policy	___	
Encryption at rest and in transit	___	
Incident response plan documented	___	
Compatible with existing systems (HMIS, EHR)	___	
Total Score	___/40	

Part 9: Frequently Asked Questions

Q: Is it legal to use AI for client documentation?

A: Yes. There is no law prohibiting the use of AI tools for documentation. The legal requirement is that you protect client privacy in the process. Using de-identified information with consumer AI tools, or using HIPAA-compliant enterprise tools with a BAA, are both legally sound approaches.

Q: What if my organization is not a HIPAA covered entity?

A: Even if HIPAA does not technically apply to your organization, you likely have privacy obligations under HMIS data standards, state privacy laws, funder requirements, or professional ethics codes. The de-identification practices in this guide represent best practice for any organization handling sensitive client information.

Q: Can I use my personal phone to access AI tools for work?

A: Check your organization's device policy. If personal devices are permitted for work, the same de-identification rules apply. Avoid saving client-related AI conversations on personal devices.

Q: What if the AI generates incorrect information?

A: This is called "hallucination" and it is common. AI may add details that were not in your notes, use the wrong tone, or make factual errors. This is why you must always read and edit the entire AI output before submitting documentation. You are the clinician; AI is the assistant.

Q: Do I need to delete my AI conversations?

A: Deleting conversations after use is best practice, especially for any conversation that involved client-related content (even de-identified). It reduces risk. Both ChatGPT and Claude allow you to delete individual conversations.

Q: What if I accidentally enter PHI into an AI tool?

A: Don't panic. Stop, delete the conversation, report it to your supervisor, and document what happened. A single accidental disclosure to a consumer AI tool, while a policy violation, is generally low risk. The important thing is to report it so your organization can learn and improve.

Q: Can I upload documents or PDFs to AI tools?

A: Do not upload any document containing client information to consumer AI tools. If a document has been fully de-identified, it may be uploaded. When in doubt, copy and paste only the specific text you need, after de-identifying it.

Q: What about AI transcription of meetings or sessions?

A: Do not upload audio or video recordings containing client information to consumer AI transcription services. Use a HIPAA-compliant transcription service if transcription of client sessions is needed, or transcribe manually and then de-identify before using AI for note formatting.

Q: Is it safe to mention the name of another agency?

A: Generally yes. Agency names are not PHI. However, if mentioning a specific agency combined with other details could identify the client (e.g., "referred to [small specialized program] for [rare condition]"), consider generalizing.

Q: What about AI features built into our EHR?

A: AI features integrated into your EHR system are typically covered by your EHR vendor's BAA and are subject to your organization's existing HIPAA compliance framework. Verify with your EHR vendor that their AI features are included in your BAA. These tools may allow direct use of PHI.

Part 10: Implementation Checklist for Leadership

Phase 1: Establish Policy (Week 1-2)

- Review this guide with your leadership team and compliance officer
- Customize the organizational AI privacy policy (Part 6) for your organization
- Identify which AI tools are approved and document their status (BAA or no BAA)
- Add AI use to your Notice of Privacy Practices (if applicable)
- Update your HIPAA risk assessment to include AI tools

Phase 2: Train Staff (Week 3-4)

- Schedule AI privacy training for all staff who will use AI tools
- Distribute the de-identification protocol and Stranger Test as quick-reference guides
- Walk through the 5-step workflow with live examples
- Practice de-identification exercises with sample notes
- Create psychological safety: emphasize that mistakes should be reported, not hidden

Phase 3: Launch and Monitor (Month 2+)

- Begin supervised AI use with de-identification protocol
- Designate an AI Champion or supervisor to answer questions and spot-check
- Conduct monthly spot-checks of documentation for any signs of PHI leakage
- Track and respond to any reported incidents
- Gather staff feedback on the workflow and refine as needed

Ongoing: Review and Improve

- Refresh training quarterly or when new staff join
- Monitor for new HIPAA/AI regulatory guidance (especially the Security Rule finalization expected mid-2026)
- Review and update approved tool list as vendors change their policies
- Share wins: track time saved and quality improvements to build organizational support
- Re-evaluate whether enterprise/BAA-covered AI tools are worth the investment

References & Resources

- HHS Office for Civil Rights. HIPAA Privacy Rule and Sharing Information Related to Mental Health. hhs.gov/hipaa
- HHS Office for Civil Rights. HIPAA Security Rule NPRM (December 2024). hhs.gov/hipaa/for-professionals/security
- HUD Exchange. HMIS Data Standards. hudexchange.info/hmis
- SAMHSA. 42 CFR Part 2: Confidentiality of Substance Use Disorder Patient Records. samhsa.gov
- NIST AI Risk Management Framework. nist.gov/artificial-intelligence
- TechSoup & TAPP Network. State of AI in Nonprofits 2025.
- HHS AI Strategy (December 2025). hhs.gov
