The AI policy your practice *should have had last year.*

How to ask your staff what AI they are already using, the conversation to have, a form to hand out, and the written policy to adopt afterward.

THIS IS A STARTING POINT, NOT LEGAL ADVICE
dX80 Group LLC is a consultation company that works alongside private medical practices, not a law firm. This document has not been reviewed by legal counsel in your state. Laws differ by state and every practice is different. Before you adopt any portion of this document, have your attorney review and approve it for use in your practice.

What's inside

- **Why it matters.** What one national survey found about unapproved AI use, and what it cannot tell you.
- **How to run it.** A dated all-staff meeting, one private conversation first, and two ways to ask.
- **The conversation.** Word for word, in four steps, with what not to say.
- **A form to hand out.** One anonymous page, ready to photocopy, if you would rather people answered on paper.
- **The policy.** Ten sections to complete with your attorney, and a signature page for your staff.

How to use it

Work through it in order. The front half is the meeting. The policy at the back is what you fill in with your attorney and adopt afterward, and the last page your staff sign is the record that they received it.

This document is provided by dX80 Group, LLC for educational and informational use only. It does not serve as legal, medical, or compliance guidance specific to your practice, and reading it does not create a consulting or professional relationship. If your practice is a HIPAA covered entity, it remains responsible for its own legal and compliance decisions. Nothing here is a guarantee of outcomes, compliance, or results. dX80 provides this free educational resource strictly for informational purposes. This is not legal advice, and it is not a draft that has been approved by your attorney. Before using it in your practice, consult your own attorney and compliance professionals. Every material statistic is tied to its original source, with derived calculations and limitations identified. If we got something wrong, tell us: results@dx80group.com.
Version 2.5, August 2026. © 2026 dX80 Group, LLC, Henderson, NV.

Why does this matter to my practice?

A December 2025 survey of 518 healthcare professionals by Wolters Kluwer Health found that 40% had encountered unauthorized AI tools in the workplace but didn't use them, and nearly 17% reported using unapproved tools themselves.[1,2]

Because that is a self-reported national sample, it cannot tell you what is true in your building, but it is enough to justify asking rather than assuming no unapproved use exists. The question is whether the wording makes truth-telling safe or dangerous. Punish the first honest answer and it will be the last one you get.

How do I actually do this?

One all-staff meeting plus a few private conversations. In this order.

1. **Pick a date for an all-staff meeting.** Everyone who touches practice information, clinical and administrative. Fifteen minutes is enough. Announce it two days ahead with Step 1.
2. **Talk to one person first.** Any time before the meeting. Next section.
3. **Choose how to run it.** The only difference is whether people answer out loud or on paper.

- **Option A, in person.** Ask the Step 3 questions out loud. Hold the patient-data question for one-to-one afterward. Fastest, if people already speak freely.
- **Option B, on paper.** Hand out the copy page at the end of this section. Anonymous return: a box on [PERSON]'s desk, back by [DATE]. Read them, then do the one-to-one conversations. Slower, better if people will not speak up in front of each other.

Steps 1 to 4 are the meeting, the same either way. The patient-data question is one-to-one in both.

First, one conversation

This conversation was written with the intention of surfacing unapproved AI use in your office. It is worded with the intent of seeking the truth from your staff rather than forcing your staff to hide.

Start with your most tech-forward staff member, any time before the group meeting. They will likely surface issues staff are less willing to raise first in a group, and they become an ally in the room instead of a defendant.

> *"You know the tools better than most people here. Before we write rules, I want your read: what's actually being used, and what should we look at first? A truthful, good-faith disclosure won't itself be the reason anyone is disciplined. I can't promise an outcome before we know the facts, but cooperating will count in your favor."*

Step 1. The announcement (send two days before)

> *"At Thursday's huddle I want to talk about AI tools, what we're already using anywhere in the office, and what we want to use. I haven't decided anything. I want to know what's real before we write any rules."*

Step 2. The opening (open with this at the huddle or your chosen meeting time)

> *"I want to talk about AI. Not because I think anyone is doing anything wrong, but because these tools are everywhere now, and I would rather we decide together how we use them than have it happen by accident."*
>
> *"So I am not asking who is using what. I am asking what is being used in this office and what it is actually helping with. 'I have seen it' is a complete answer. What I want is to get the useful ones set up properly instead of hidden."*
>
> *"One thing I do need individually. If anything patient-related has ever gone into one of these tools, even a name or a partial detail, come and find me and tell me one to one. I need to know so I can handle it properly, and I would far rather hear it from you than find it later."*

Do not promise blanket amnesty in the room. You cannot deliver it and your staff know it, which is why promising it costs you the credibility the rest of this depends on. What is true, and what you can say if someone asks: a truthful, good-faith disclosure will not, by itself, be the reason anyone is disciplined. Consequences for the underlying conduct are handled case by case under your existing policies, considering things like intent, whether it kept happening, and cooperation. Say that when you are asked. Do not lead with it.

Step 3. The four questions

Ask these out loud at the huddle, or put them on paper as in Option B and hand them out with a named return point and a date. The ready-to-copy version is the page at the end of this section.

- "What tools have you used or seen used here, even once?" (Write every answer down without commentary. Your face is the policy right now.)
- "What were you trying to get done when you reached for it?" (The task list is your real gold: it is the map of where your workflows are failing your team.)
- "Did anything patient-related ever go into one of them, even names or partial details?" (Ask it flat and calm. You need this answer for your own risk assessment. React with gratitude, not alarm, or you will never hear a true answer again.)
- "If you could wave a wand, what would you want AI to take off your plate?" (Ends the meeting forward-looking, and tells you what to vet first.)

Hold the third question back for a one-to-one conversation, in both options. If the answer is yes you may have an obligation to look into it, and you cannot look into something raised anonymously or in a group where nobody will claim it.

What not to say

- Do not say "who authorized this?" Nobody did. That is the point of the meeting.
- Do not name or look at a specific person while asking about tools.
- Do not promise "we'll approve everything." You won't. Promise a fast, honest review.
- Do not end without a date: "Approved list published by [DATE]. Until then, nothing patient-related goes into anything new."

Step 4. After the meeting

- Same week: write down every tool named, into the Approved AI Tools List template, status "under review."
- If PHI may have entered an unapproved tool: contact your attorney or compliance advisor about assessment obligations. Do it calmly and promptly. This is why you asked.
- Within two weeks: publish the first approved list, even if it is short. A fast, honest "not yet" beats a slow silence.
- Adopt the written policy that follows at a follow-up meeting, with signatures.

The next page is the form for Option B. Write in the return name and date, then photocopy it. Do not add a name field, and note that the patient-data question is deliberately not on it — that one is a conversation with you.

AI tools in our office

We are making a list of the AI tools already in use here, so the useful ones can be approved instead of hidden. No name needed.

1. What AI tools have you used or seen used here, even once?

2. What were you trying to get done when you reached for it?

3. If you could wave a wand, what would you want AI to take off your plate?

Return this to _____

by _____

[PRACTICE NAME] AI Use Policy

Policy control — complete on adoption: Effective date _____ · Adopted by _____ · Policy owner _____ · Version _____ · Next review _____

1. Purpose and scope

[PRACTICE NAME] permits AI tools only as reviewed and approved under this policy when the use involves practice information — PHI, confidential or internal business information, patient communications, clinical content, employment or personnel information, credentials, or financial data — and prohibits all other AI use of that information. Public, non-confidential information (for example, already-published material) may be used only as the practice authorizes and subject to accuracy and intellectual-property requirements. This policy applies to every owner, clinician, employee, contractor, and student, on any device, personal or practice-owned, whenever practice information is involved.

This policy supplements and does not replace the practice's existing HIPAA Privacy, Security, sanctions, and breach-notification policies; where they overlap, the more protective requirement governs, and nothing in this policy waives any obligation under HIPAA or the practice's other policies. [ATTORNEY/PRACTICE: confirm precedence for our practice.]

2. Definitions

- **AI tool:** any software that generates, summarizes, transcribes, drafts, or decides using machine learning or a language model. This includes chatbots, ambient scribes, dictation with AI features, browser extensions, and AI features inside existing products (EHR, email, office suites).
- **PHI: individually identifiable health information protected by HIPAA when created, received, maintained, or transmitted by a covered entity or business associate. It may include names, dates, contact information, images, recordings, medical details, and combinations of information that identify or could reasonably identify an individual.**
- **Approved tool:** a tool listed on the practice's Approved AI Tools List, used within the limits written there.
- **Business associate agreement (BAA):** the written contract HIPAA requires before a vendor creates, receives, maintains, or transmits PHI on the practice's behalf.[3]

3. The core rules

- **Rule 1.** No PHI enters any AI tool that is not on the Approved AI Tools List with "PHI permitted" marked yes and a signed BAA on file.
- **Rule 2.** Every AI-generated clinical note, letter, or message is reviewed and approved by the responsible clinician before it is signed, sent, or filed. The clinician must review the note for accuracy, completeness, and clinical appropriateness before approving or signing it, and remains responsible for the professional, documentation, billing, and organizational obligations that apply to the final note.[4]

- **Rule 3.** No recording of patient encounters without the notice and consent procedure in Section 6.
- **Rule 4.** AI output never substitutes for clinical judgment. It drafts; it does not decide.
- **Rule 5.** Staff complete the training in Section 7 before receiving access to any approved tool. Training first, tools second.

4. Approved tools and the approval process

The practice maintains an Approved AI Tools List (template provided in this library) naming each approved tool, what it may be used for, whether PHI is permitted, the BAA date, and the review date. [ROLE, e.g., PRACTICE OWNER or PRIVACY OFFICER] owns the list. To request a new tool, staff submit the request to the list owner, who vets it using the practice's vendor checklist (BAA and Vendor-PHI Checklist, and the AI Scribe Vetting Checklist for documentation tools) before any use begins.[5]

5. PHI and HIPAA requirements

- A vendor whose tool creates, receives, maintains, or transmits PHI on our behalf is a business associate, and a BAA meeting 45 CFR 164.504(e) must be signed before any PHI flows (see HHS sample provisions, cited below).[6]
- Where the minimum necessary standard applies — and as a prudent data-governance default even where a technical exception may exist — an AI tool receives only the PHI its approved task requires (45 CFR 164.502(b)).[7]
- Adding or materially changing an AI tool that creates, receives, maintains, or transmits ePHI requires the practice to assess and document the resulting risks as part of its ongoing Security Rule risk-analysis and risk-management process, and to update safeguards or the documented analysis as appropriate (45 CFR 164.308(a)(1)(ii)(A) and OCR risk analysis guidance).[8]
- Removing patient names alone does not de-identify PHI. De-identification requires the Safe Harbor removal of 18 identifier categories or expert determination (45 CFR 164.514).[9]
- Classify practice information before it goes near AI. A simple scheme works — public, internal, confidential, and restricted. PHI and any identifiable patient information is restricted: it never enters an AI tool that lacks the "PHI permitted" approval and a signed BAA, and never a personal or consumer account.
- **State law may be stricter than HIPAA.** State privacy, medical-records, breach-notification, and recording laws can impose requirements beyond HIPAA; confirm your state's rules, and any telehealth patient's state, with counsel.
- **Substance-use and behavioral-health records carry extra rules.** If the practice handles 42 CFR Part 2 substance-use-disorder records, or is subject to state behavioral-health or minor-consent confidentiality laws, additional restrictions apply before any such information reaches an AI tool — confirm with counsel.

6. Patient notice and consent for recording

Before any encounter is recorded for ambient documentation: [DESCRIBE PROCEDURE, e.g., verbal notice at the start of each visit plus posted signage, with consent documented in the

record and an opt-out honored without affecting care]. Several states require all-party consent to record a conversation; this practice treats all-party consent as the default everywhere. [ATTORNEY: confirm requirements for OUR STATE and any state our telehealth patients are located in.]

7. Training

Before access, each user completes: how the approved tool works and its known failure modes (for scribes: omission errors are the most frequent error type found in validation research)[10]; what PHI may and may not enter it; the review-before-signing requirement; and how to report a suspected error or breach. Refreshed [ANNUALLY].

8. Good-faith incident reporting and non-retaliation

Anyone who suspects PHI entered an unapproved tool, or that an AI error reached a patient record, reports it to [PRIVACY OFFICER / OWNER] the same day. The practice prohibits retaliation against an individual because the individual made a good-faith report, requested guidance, or took part in an investigation; prohibited retaliation includes termination, demotion, suspension, reduced compensation, reassignment, or any other adverse action taken because of a good-faith report. That protection does not shield the underlying conduct: the practice may still evaluate it under its existing privacy, security, sanctions, and employment policies, considering intent, recurrence, and cooperation, and a knowingly false or malicious report is itself subject to discipline. Nothing in this section alters the at-will employment relationship or creates any contractual right; only a written agreement signed by an authorized representative of the practice can do that. The goal of this policy is a managed practice, not a scared one. [ATTORNEY: confirm this non-retaliation and at-will language for our state and employment posture.]

A report under this section starts, and does not replace, the practice's breach-assessment process: promptly assess the nature and scope, contain it where practicable, and determine whether notification is required under 45 CFR 164.400–414 and any applicable state breach-notification law (for example, in Nevada, NRS Chapter 603A). Notify affected individuals, HHS, and any others as that assessment and those laws require.

9. Personal devices and personal accounts

Practice information never enters personal AI accounts (personal chatbot logins, free consumer apps) regardless of device. Approved tools are used only under practice-controlled accounts, so the practice can enforce settings, retention, and the BAA.

10. Review

This policy is reviewed on a set interval (for example, every 6 or 12 months, as the practice decides), and immediately upon: a new tool approval, a reported incident, or a change in law or guidance. Note: HHS has proposed updates to the HIPAA Security Rule (proposed rule published

January 6, 2025); the current Security Rule remains in effect while rulemaking proceeds.[11,12] Your attorney should confirm current status at each review.

Acknowledgment (each staff member signs)

I have read and understood the [PRACTICE NAME] AI Use Policy. I agree to use only approved AI tools, within their approved limits, and to report suspected incidents the same day I become aware of them.

Name: _____

Role: _____

Signature: _____

Date: _____

REFERENCES

Check our work.

Numbered in the order they are cited.

1. Wolters Kluwer Health. Wolters Kluwer survey finds broad presence of unsanctioned AI tools in hospitals and health systems. January 22, 2026. Survey of 518 healthcare professionals conducted December 2025. Accessed August 30, 2026. https://www.wolterskluwer.com/en/news/wolters-kluwer-survey-finds-broad-presence-of-unsanctioned-ai-tools-in-hospitals-and-health-systems

2. Wolters Kluwer Health. Shadow AI: providers are using unapproved tools to improve workflow. January 22, 2026. Source of the 17% figure: "Across all respondents, 17% admitted to using unauthorized AI tools in the workplace." Accessed August 30, 2026. https://www.wolterskluwer.com/en/expert-insights/shadow-ai-providers-are-using-unapproved-tools-to-improve-workflow

3. US Department of Health and Human Services. Business associates. Accessed August 30, 2026. https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/business-associates

4. American Medical Association. Augmented Intelligence Development, Deployment, and Use in Health Care. Board of Trustees Report. November 2024. Accessed August 30, 2026. https://www.ama-assn.org/system/files/ama-ai-principles.pdf

5. Lozovatsky M, Thomas M. Governance for augmented intelligence: establish a governance framework to implement, manage, and scale AI solutions. AMA STEPS Forward. 2025. Accessed August 30, 2026. https://edhub.ama-assn.org/steps-forward/module/2833560

6. US Department of Health and Human Services. Business associate contracts: sample business associate agreement provisions (45 CFR §164.504(e)). January 25, 2013. Accessed August 30, 2026. https://www.hhs.gov/hipaa/for-professionals/covered-entities/sample-business-associate-agreement-provisions

7. US Department of Health and Human Services. Minimum necessary requirement. Accessed August 30, 2026. https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/minimum-necessary-requirement

8. US Department of Health and Human Services, Office for Civil Rights. Guidance on risk analysis. Accessed August 30, 2026. https://www.hhs.gov/hipaa/for-professionals/security/guidance/guidance-risk-analysis

9. US Department of Health and Human Services. Guidance regarding methods for de-identification of protected health information in accordance with the HIPAA Privacy Rule (45 CFR §164.514). Accessed August 30, 2026. https://www.hhs.gov/hipaa/for-professionals/special-topics/de-identification

10. Biro J, Handley JL, Cobb NK, et al. Accuracy and safety of AI-enabled scribe technology: instrument validation study. J Med Internet Res. 2025;27:e64993. doi:10.2196/64993

11. US Department of Health and Human Services. HIPAA Security Rule NPRM. Accessed August 30, 2026. https://www.hhs.gov/hipaa/for-professionals/security/hipaa-security-rule-nprm

12. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. US Dept of Commerce; January 2023. Accessed August 30, 2026. https://www.nist.gov/itl/ai-risk-management-framework

This document is provided by dX80 Group, LLC for educational and informational use only. It does not serve as legal, medical, or compliance guidance specific to your practice, and reading it does not create a consulting or professional relationship. If your practice is a HIPAA covered entity, it remains responsible for its own legal and compliance decisions. Nothing here is a guarantee of outcomes, compliance, or results. dX80 provides this free educational resource strictly for informational purposes. This is not legal advice, and it is not a draft that has been approved by your attorney. Before using it in your practice, consult your own attorney and compliance professionals. Every material statistic is tied to its original source, with derived calculations and limitations identified. If we got something wrong, tell us: results@dx80group.com.
Version 2.5, August 2026. © 2026 dX80 Group, LLC, Henderson, NV.