TIER 1 - FOUNDATIONS * V1.0 -- MAY 2026

PRIVACY HYGIENE FOR AI

The full posture for personal AI use. Account separation, what to never paste, the gap between what each tool says it does with your data and what actually happens, and the pre-mortem for the moment you wish you hadn't shared something.

BY Alex Jahn / Agent Logic

v1.0 -- May 2026 **EDITION**

Anyone who's been using AI casually for a while and wants to clean up their privacy posture before it bites them -- or before they share something sensitive on a system they don't actually understand **AUDIENCE**

15-20 minutes **READ TIME**

Free. Forever. **COST**

Prepared by Agent Logic / alexanderjahn79@icloud.com / theaiguywi.com

CONTENTS

What's in here

SECTION 1

The data you didn't realize you gave

Most people are running on default settings

When you signed up for ChatGPT (or Claude, or Gemini, or Copilot), you probably clicked through the terms-of-service in 4 seconds and accepted the default privacy settings. Most people do. The defaults are usually the company's preference, which is usually more data collection rather than less.

A year of use later, the typical AI account contains: every conversation you've had with the model, including the personal stuff (your kids' names, your client situations, your money decisions, your health questions), with most of it possibly being used as training data, possibly being reviewed by humans for "safety," possibly being retained indefinitely, and possibly being subject to a future policy change that retroactively changes how it's all handled.

That's the worst-case version. The best-case version is "actually they delete most of it after 30 days and nobody looks at it." The honest version is somewhere in between, varies by tool, and is harder to pin down than the marketing pages suggest.

This module is the cleanup. The full privacy posture for personal AI use -- assuming the worst-case more than the best-case, because the asymmetry of being wrong runs in only one direction.

What you'll have by page 13

By the end of this primer:

- The **assume-l**
- A **3-tier sep**
  goes where, and why.
- The **never-pas**
- The **per-tool p**
  about your data, in plain language.
- A **15-minute**
- The **pre-morte**
- The **honest lir**

> ***Treat anything you say to a chat tool like a phone call you can't take back. The privacy posture in this module is the version of carefulness that lets you keep using AI without worrying about what's downstream.***

SECTION 2

The assume-leaks rule

The foundational stance

> **The assume-leaks rule:**

For privacy decisions, assume that anything you put into a chat tool may eventually be:

- Used as training data for future models
- Reviewed by a human (employee, contractor, or auditor)
- Subject to a future policy change that retroactively expands what's done with it
- Accessible to anyone with admin or legal access to the company's servers
- Visible in a future data breach

You're not assuming the company is acting in bad faith. You're assuming that the worst-case version of any of these things is possible -- and adjusting what you share accordingly.

The opposite stance -- "I'm sure they wouldn't actually do that" -- is what most people implicitly use, and it's the stance that gets people hurt when policies change, breaches happen, or "what's safe today" turns out not to be what's safe tomorrow.

Why this isn't paranoia

The actual evidence:

- **Multiple major AI tools have had bugs that exposed one user's data to another.** ChatGPT had this in 2023. Subsequent occasional incidents have surfaced for other tools. None of these were intentional; all of them happened.

- **Privacy policies change.** A tool can have a "we don't train on your data" policy today and a "we do train on your data unless you opt out" policy in 18 months. Existing data is often subject to the new policy when this happens.
- **Acquisitions change everything.** If the company you trusted with your data gets acquired, the new owner can rewrite the rules. This has happened in the broader software industry many times.
- **Employee access varies.** "We don't review your conversations" is sometimes "we don't ROUTINELY review your conversations" with significant exceptions for safety, abuse, and quality testing.

The assume-leaks rule isn't accusing anyone. It's accepting that you don't have control over the variables, and adjusting on the inputs you DO control.

SECTION 3

The 3-tier separation framework

Three categories of content; three different homes

> **The 3-tier separation:**

Tier A -- Work content. Goes in your work-issued account or the work-tier of a tool. Subject to your employer's data-handling policies. Don't put personal stuff here; if you leave the job, you may lose access to your own conversations.

Tier B -- Personal content. Goes in your personal account. Lower-stakes day-to-day questions, life-admin tasks, learning conversations. Most of your daily AI use lives here.

Tier C -- Sensitive personal content. Doesn't go in a chat tool at all without specific extra precautions. Health, money, family, relationships, legal -- the topics where being wrong about the privacy posture has real consequences.

The mistake most people make is putting all three tiers in the same account, on default settings. The fix is mechanical: separate accounts, intentional decisions about which tier each conversation belongs to.

Setup specifics

Tier A -- work account:

- Use the work-issued AI tool if there is one (Microsoft 365 Copilot, ChatGPT Enterprise, Claude for Work). These have stronger data-handling guarantees built into the contract your employer signed.
- If your employer doesn't issue a tool but allows AI use, sign in with your work email -- this signals intent and creates a paper trail you can defend.
- Don't paste personal context into the work account. Your employer's IT may have access.

Tier B -- personal account:

- Sign up with your personal email. Different account from work.
- Pick the tool with the privacy posture you trust most for your personal use (see Section 5 for per-tool comparison). For most people, this is Claude.
- Turn OFF training-data sharing if your tool has a toggle for it.
- Use this for daily-learning, life-admin, the bulk of your AI use.

Tier C -- sensitive content:

- For most sensitive topics, don't use a chat tool at all without extra precautions. The doctor / lawyer / accountant is a better fit.
- If you must use a chat tool for sensitive content (the topic was real, the question urgent, the alternative was worse): use a paid tier, training off, in a deliberate "incognito" or no-save mode if available, and treat the conversation as needing extra scrub.
- Never use the same tool for Tier C and casual Tier B in the same session -- context bleeds.

This sounds elaborate. The setup is a one-time investment of about 20 minutes. After that, the muscle memory ("which account am I in?") is automatic.

SECTION 4

The never-paste list

Data types that should never enter a chat tool, period

> **The never-paste list:**

- **Social Security numbers, tax ID numbers, full date of birth combined with full name.**
- **Bank account numbers, routing numbers, full credit card numbers.**
- **Passwords. Even temporarily. Even "I'll just paste it to ask the AI to fix the format."**
- **API keys, access tokens, security credentials of any kind.**
- **Specific addresses of your home or your kids' schools.** General location ("Fond du Lac, WI") is fine; the actual house number is not.
- **Other people's full identifying information** -- your spouse's medical record, your client's full account details, your coworker's name + situation. They didn't consent.
- **Active legal documents in disputes you're a party to.** Lawyer-only territory.
- **Therapy notes, prescription records, current diagnosis names** combined with your full identity. Health information has special legal status; treat it that way.

Memorize this list. Internalize it. The rule isn't "be reasonable" -- it's that these specific data types are bright lines.

Why each one

- **SSN / tax ID / DOB+name** -- identity-theft vectors. Once leaked, can't be revoked easily.
- **Bank / card numbers** -- fraud vectors. Same problem.
- **Passwords** -- account-takeover vectors. AI tools sometimes log entire prompts; a password in a prompt is a password in a log.
- **API keys / tokens** -- service-takeover vectors. Same as passwords but for systems instead of accounts.
- **Specific addresses / kids' schools** -- physical safety. AI logs are searchable in ways your home address shouldn't be.
- **Other people's data** -- consent + ethics. Even if you think it's fine, they didn't agree.
- **Active legal documents** -- privilege + strategy. Lawyer's job, not chat tool's.
- **Health information** -- legal protection (HIPAA-adjacent for some) + future-policy risk.

The list is short on purpose. There's plenty you CAN paste. These eight categories are the ones where being wrong is irreversible.

SECTION 5

Per-tool privacy snapshot (mid-2026)

What each major tool says about your data. Plain language. The space changes; verify the live policy when stakes are high.

ChatGPT (OpenAI) -- Free tier may use your conversations as training data by default. Plus and Pro tiers have a setting to turn this off (Settings -> Data Controls -> "Improve the model for everyone"). Enterprise / Team tiers don't train by default. There's a "Temporary Chat" mode for sessions you don't want saved. Practical implication: if you're on the free tier, treat as if training is on. Paid + toggle-off is reasonable for personal use.

Claude (Anthropic) -- Does NOT train on consumer conversations (free or Pro) by default. This is a real differentiator among the four. Conversations may be reviewed for safety and policy violations, not for model training. Practical implication: among the four, Claude has the cleanest default privacy posture for individual users. Still subject to assume-leaks, but the policy is the most consumer-friendly.

Gemini (Google) -- Conversations may be reviewed by humans for quality and may be used to improve services. Activity controls let you turn off conversation history, which weakens the assistant's continuity. Practical implication: assume Google sees the data. Use Gemini for low-sensitivity tasks; use a different tool for sensitive personal content.

Copilot (Microsoft) -- Consumer Copilot has commercial-data-protection on by default for signed-in users -- better than the typical consumer norm. Microsoft 365 Copilot (the paid Office version) has enterprise-grade controls. Practical implication: Copilot is reasonable for personal use, especially if you already have a Microsoft 365 subscription.

> **The simple rule for tool selection by sensitivity:**

- **Daily / non-sensitive** -> any tool. Convenience wins.
- **Personal but sensitive** (money, family, health, decisions) -> Claude (best default), or paid tier of others with training-off toggle confirmed.
- **Other people's data** -> no consumer chat tool. Use a business-tier with proper data-handling contracts.

- **Legally privileged** -> no chat tool. Period.

SECTION 6

The 15-minute privacy audit

A one-time cleanup of your existing accounts. Do this once, then re-do every 6 months as policies change.

Audit checklist

For each AI tool you use:

> **The 15-minute audit:**

1. **Open Settings -> Data Controls** (or Privacy, or whatever the tool calls it).
2. **Toggle OFF training-data sharing** if the tool offers this and you're on a tier where it works.
3. **Look at the saved memories or conversations list.** Read what's there. Delete anything that's wrong, outdated, or shouldn't have been saved.
4. **Check for "anyone can see your conversations" or sharing settings.** Turn off public sharing unless you specifically use it.
5. **Set a policy for retention.** If the tool offers automatic deletion (e.g., "delete chats after 30 days"), turn it on for personal accounts unless you have a reason not to.
6. **Verify the email / phone associated with the account** is the right tier (work vs personal). If you signed up with work email by accident, decide whether to migrate or close.

For each session going forward:

- Check which account you're in BEFORE pasting anything sensitive.
- Use Temporary Chat / private mode for genuinely sensitive single-conversation work.
- Never re-paste the same sensitive content into multiple tools "just to check the answer." That doubles the privacy exposure.

That's the audit. 15 minutes, once, plus the per-session muscle memory.

SECTION 7

The pre-mortem move

The single best privacy habit

Before you paste anything you're not 100% sure about, stop and ask:

> **The pre-mortem question:**

"If this conversation showed up screenshotted in an email tomorrow, which version of me would be embarrassed -- and is there a version of this question I can ask without that risk?"

Three possible outcomes:

- **No embarrassment risk.** Paste freely.
- **Some embarrassment risk, but the question can be rephrased.** Rephrase to ask the structural version. ("How do I think about negotiating with a coworker who took credit for my work" instead of "Help me write an email to Ryan about Wednesday's meeting." The first protects you; the second names a colleague.)
- **Real embarrassment risk and the question can't be rephrased without losing the point.** Don't paste. Talk to a human -- friend, lawyer, therapist, professional in the relevant field.

The pre-mortem takes 5 seconds. It's the single highest-ROI privacy habit you can build, and it works across every tool, every tier, every situation.

> ***The cost of pausing 5 seconds to ask "what's the worst case?" is nothing. The cost of not pausing, when the worst case actually happens, is unbounded.***

SECTION 8

When privacy posture isn't enough

Three categories where good hygiene still isn't sufficient:

- **Active legal exposure.** If you're a defendant in a lawsuit, under investigation, going through divorce or custody -- your AI conversations are potentially discoverable. The privacy settings don't protect you from a subpoena. Don't use chat tools to discuss anything related to the matter; talk to your attorney.
- **Sensitive employment situations.** Filing a complaint about your employer, planning to leave, discussing a coworker -- if your employer's IT has any reach into your accounts (especially if you've used a work email or work device), assume readable. Use a personal device, personal account, neither of which has been used for work tasks before.
- **Active health crises.** If something is genuinely dangerous and time-sensitive, don't ask AI -- call a doctor or 988 or 911. AI can be useful for understanding a chronic condition; it should not be the first stop in an emergency.

Within those limits, the posture in this module covers normal personal AI use thoroughly enough that you can stop worrying about background privacy risk and focus on the actual conversations.

SECTION 9

Where to go from here

You have the privacy posture. Two more modules in the Tier 1 expansion:

- **AI for travel and everyday research** -- vacation planning, hotel comparison, road-trip regulations. Where multi-factor structuring shines and where current local data falls apart.

After that: recovering when AI is wrong. Tier 1 closes at 18 modules.

Get the next module the day it drops: theaiguywi.com/training

One email per release. No drip. No spam. Opt out anytime.

If you want this same privacy posture installed across a small business -- the team trained once on the assume-leaks rule, the never-paste list adopted shop-wide, the audit done across every employee account -- that's the consulting offer.

Reach out: alexanderjahn79@icloud.com

A short call. Honest scope. We figure out together if it's a fit.

Closing -- the lock-in line

Privacy isn't paranoia. It's accepting that you don't control how the company handles your data over time, and adjusting on the inputs you DO control. The 3-tier separation, the never-paste list, the assume-leaks rule, and the pre-mortem question together form a posture that lets you keep using AI without quietly worrying about what's downstream.

8

Eight items on the never-paste list.

Memorize them. They're the bright lines that protect you from the irreversible mistakes -- identity theft, account takeover, legal exposure, broken trust with people whose data isn't yours to share.

You have the posture. Two more modules in this batch.

-- Alex

Agent Logic -- *theaiguyw*

Fond du Lac, WI. This is module 16 of 18 in Tier 1 (Personal).