TIER 2 - PROFESSIONAL * V1.0 -- MAY 2026

PRIVACY AT WORK

The workplace data classifications, consumer vs enterprise tier differences, and the specific policy decisions every team needs to make before the first prompt goes out the door.

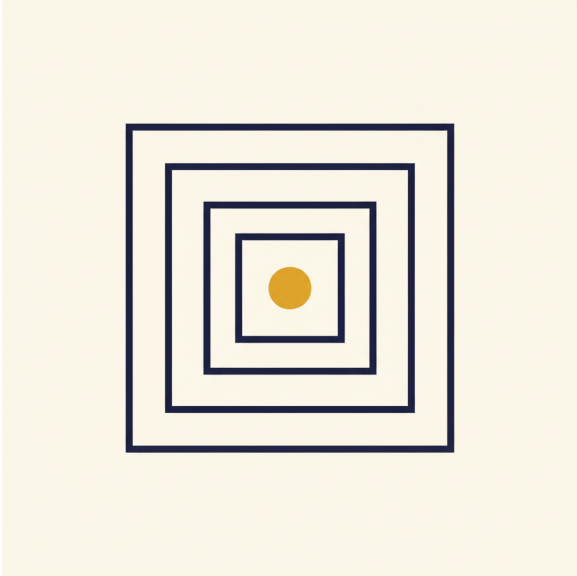

BY Alex Jahn / Agent Logic

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

Anyone using AI for work at a company that has customers, employees, vendors, or contracts -- and the operator who has to set the policy **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 other half of the red line

What you can't afford to leak from the company

Tier 1 covered the personal half of privacy -- your own SSN, your family's medical info, your kid's school records. The cost of leaking that data is real but mostly contained: it's your data, your decision, your cost.

This module is the other half. **What you**
bills, or the customers who trust you with their information. The cost of leaking THIS data is bigger and the failure modes are more public:

- A customer's tax data ends up in a model's training corpus. They find out. You're now in a months-long disclosure conversation, possibly a lawsuit, definitely a relationship ender.
- An employee's salary range gets pasted into ChatGPT to "help draft a fair raise letter." It surfaces six months later in a class-action discovery. HR has questions.
- A confidential vendor pricing sheet ends up logged on a cloud provider you don't have a contract with. Procurement just lost three months of negotiating leverage.
- Your trade-secret formula or proprietary process gets pasted into a chat window. Your competitive advantage is now in a model's training data forever.

These aren't theoretical. Each one has happened to operators I know in the last 18 months. The damage was real. The fix in every case was *a clear po*
the door.

This module is that policy.

What this module covers

By the end of this primer:

- You'll know the **four class**
- You'll understand the difference between **consume** what each actually does for your data.
- You'll have a **decision**

- You'll have a **redaction**
- You'll have a working **1-page AI**
- You'll know what to say in a **SOC 2 co**

It's the longest module in Tier 2 because it's the most operationally dense. If you're an operator about to roll AI out across a team, this is mandatory. If you're an employee using AI at work, the first half is mandatory; the policy template is for your boss.

> ***"We didn't think to write a policy" is the most expensive sentence small-business operators have ever said about AI. The policy takes an afternoon. The cleanup if you don't have one takes months.***

SECTION 2

The four classes of workplace data

Public, internal, confidential, regulated

Almost every data-classification framework -- corporate, government, regulatory -- sorts business data into four classes. The names vary; the classes are the same. They map directly to AI policy decisions.

The four classes -- what each is and what it implies:

1. **Public** -- Already published. Marketing materials, public website content, press releases, anything intentionally externally visible. *AI implica*
2. **Internal** -- Visible to anyone in the company, but not to outsiders. Employee handbook, internal process docs, team-only Slack channels, project plans not yet announced. *AI implica*
   consumer tiers OK with care. Avoid pasting into models that train on inputs.
3. **Confidential** -- Visible only to a subset of the company. Salary data, contract terms, strategic plans, M&A discussions, performance reviews. *AI implica*
   contracts only.

4. **Regulated** -- Subject to specific laws. HIPAA-protected health info, GDPR-personal data, SOX-protected financial info, attorney-client communications, anything under an NDA. *AI implication:* enterprise tier with the right DPA, or don't use AI at all.

That's it. Four classes. Almost every data-handling decision at a small business reduces to "*which class is this data in?*" and "*what's the right tier for this class?*"

Examples per class -- typical small business

To make this concrete, here's how the four classes look at a typical small operating business:

Public data at a small business:

- The pricing on your website
- Marketing emails
- Customer testimonials they've approved for use
- Job postings
- Anything in a press release or blog post

Internal data at a small business:

- Your employee handbook
- Your standard operating procedures
- Project plans before launch
- Team meeting notes (without confidential decisions)
- Vendor lists

Confidential data at a small business:

- Employee salaries, comp structures, performance reviews
- Customer contracts (especially terms or pricing)
- Strategic plans (expansion, hiring, M&A talks)
- Financial details (cash position, runway, profit margins)
- Vendor contracts and pricing sheets

Regulated data at a small business:

- Customer financial info (account numbers, full SSNs, full birthdates)
- Health information (HIPAA -- even mentioning a customer's medical condition triggers it)
- Anything under a signed NDA

- Attorney-client communications
- Specific industry-regulated data (insurance, financial services, healthcare)

If you can't classify a piece of data, default to one class higher than your gut says. The cost of over-protecting is small. The cost of under-protecting is large.

SECTION 3

Consumer vs enterprise tier -- what each actual

The plain-English version

Most people have no idea what they're agreeing to when they sign up for a consumer AI tool. The fine print matters more than people realize. Two bright lines:

Consumer tier (the $20-30/month subscriptions -- ChatGPT Plus, Claude Pro, Gemini Advanced):

- **Training opt-in by default.** Your prompts can be used to improve the model unless you find the setting and opt out. Most users haven't found the setting.
- **Logged for safety review.** Every prompt is retained, typically for ~30 days, and reviewable by the provider's safety team for abuse detection.
- **Standard customer support.** No SLAs, no compliance review available, no data-processing addendum.
- **Suitable for:** Public data freely. Internal data with care. Personal use. NOT confidential or regulated.

Enterprise tier (the contractual versions -- ChatGPT Enterprise, Claude for Enterprise, Gemini for Workspace):

- **No-training contractual default.** Your prompts are explicitly excluded from training. Written into the contract.
- **Stronger logging boundaries.** Logs are typically encrypted, retention is shorter, and access is restricted to incident response only.
- **Data-processing addendum (DPA) available.** A real legal document spelling out what they do with your data and what their obligations are.
- **SOC 2 Type II report typically available.** The third-party audit confirming they actually do what they say.

- **Cost: ~3-5× per seat compared to consumer.** $30-100/seat/month at the lower end of enterprise; $100-300/seat at the higher end.
- **Suitable for:** All four classes, with the caveat that regulated data still requires you to verify the specific compliance posture matches your obligation.

> **The 60-second mental model:**

- *Consumer tier* = "you're paying for software, your data is part of how they pay for it."
- *Enterprise tier* = "you're paying enough that your data isn't part of the deal."

If you can't afford enterprise for some classes of data, you have to either redact the data, switch tools for those tasks, or do them without AI.

What's "Business" or "Team" tier?

Most providers offer a middle tier (Team, Business, Pro for Teams) that sits between consumer and enterprise. It usually includes:

- No-training default (good).
- Centralized billing (good).
- Some admin controls (good).
- BUT no DPA, no SOC 2, no compliance review (limited).

For most small businesses, Team tier is fine for *confidential* data. For *regulated*

The line moves with company size and industry. A 10-person consultancy can probably get away with Team tier for most things. A 10-person healthcare provider absolutely cannot.

SECTION 4

Which tier for which class of data

The decision matrix

Pin this on the wall. It's the single most-referenced artifact your team will use for the next year:

The data class -> AI tier decision matrix:

| Data class | Acceptable AI tiers | Notes | |-----------|--------------------|-------| | **Public** | A
team / enterprise) | Already public. Fair game. | | **Internal** |
+ redaction. | Most daily AI work happens here. | | **Confiden**
contracts, strategic plans. Need no-training contract. | | **Regulate**
compliance. Or don't use AI. | HIPAA, GDPR, SOX, NDA-bound. The lawyer needs to sign off. |

That's the whole policy in one table. The rest of this module is how to live by it.

What "with care" means for internal data on consumer tier

If you're a small team running on consumer tier and most of your work is internal data, you can stay in compliance with three habits:

1. **Opt out of training.** ChatGPT, Claude, Gemini -- all have settings to disable using your prompts for training. Find them and turn them off. (This is a one-time setting, takes 30 seconds, almost nobody does it.)
2. **Redact identifiers.** Names, account numbers, addresses -- strip them before pasting. The actual data the model needs to do the task is usually preservable without identifying info.
3. **Don't paste documents wholesale.** Summarize first, then paste. The full contract / full HR doc / full project plan rarely needs to enter the chat window. The relevant 200 words usually does.

Three habits. Five extra seconds per prompt. Reduces your exposure dramatically without requiring an enterprise upgrade.

When to upgrade to enterprise tier

The trigger isn't a fixed company size. It's a fixed *data-hand*
when:

- You handle confidential data routinely and your current consumer-tier exposure is real.
- You handle any regulated data and your provider's enterprise tier offers an applicable DPA.
- A customer or partner contractually requires it (this happens -- I've seen contracts that say "your AI vendors must be SOC 2 Type II").
- Your team is large enough that informal "be careful what you paste" doesn't scale (usually past 10-15 people).

Don't upgrade just because the marketing says "enterprise-grade." Upgrade because your data-handling reality requires it.

SECTION 5

The redaction workflow

Three minutes to clean a prompt before pasting

The single most-leveraged habit in workplace AI use. Most data-leak failures happen because someone pasted the original document instead of a redacted version. Three-step workflow that takes about 3 minutes per use:

> **The minimum-viable redaction checklist:**

1. Identify the load-bearing fact. What does the model actually need to do the task? Usually it's the structure, the situation, the context -- not the names, numbers, or identifying details.

2. Replace identifiers with placeholders or roles.

- Names -> roles:
- Account numbers -> placeholders:
- Specific dollar figures -> ranges:
- Specific addresses -> general locations: *residential property"*
- Dates that uniquely identify -> relative:

3. Re-read the redacted version. Could a reasonable outsider re-identify the customer or employee from this prompt? If yes, redact more. If no, you're safe to paste.

That's it. Three minutes. Reduces 95% of accidental data leaks.

When redaction isn't enough

For some tasks, the data IS the identifying detail. You can't do "draft this customer's tax filing" with the customer's name redacted, because the filing requires the name. For tasks like that, redaction isn't the right tool -- you need the right tier (enterprise + DPA) or you need to do the task without AI.

Pattern:
identifying detail IS the task.

Synthetic substitution -- the advanced move

Sometimes you can substitute realistic fake data instead of redacting. Useful when the task needs the *shape* of the data but not the actual values.

- Generating sample contracts? Use a fake customer name, fake dollar figures, fake project description that mimics your real one.
- Stress-testing a customer email? Use a synthesized customer scenario that captures the situation without being a real person.
- Training your team on AI? Use synthetic case studies, not real customer files.

The model's output is just as useful -- and you've moved zero real data through the chat window. Worth the 5 extra minutes for high-frequency template tasks.

SECTION 6

A 1-page AI usage policy that actually works

The minimum policy a small team needs

Most small businesses don't have an AI usage policy. Most should. Here's the minimum-viable policy that fits on one page and answers the questions your team will actually have.

Adapt to your situation; don't copy verbatim:

> **[Company Name] -- AI Usage Policy v1.0**

1. Approved tools. We use [vendor name] [tier -- Team / Enterprise]. No other AI tools without prior approval. Personal accounts may not be used for company work.

2. What you may paste:

- Public data: freely.
- Internal data: yes, with judgment. Don't paste full documents -- summarize and paste the relevant excerpt.
- Confidential data (salaries, contracts, strategic plans): only with [Owner] approval, and only on [Enterprise tier].

- Regulated data (customer financial info, health info, NDA-bound material): never via [Team / consumer tier]. Contact [Owner] before using AI on regulated data even at enterprise tier.

3. Redaction. When in doubt, redact. Replace customer/employee names with roles, account numbers with placeholders, specific dollar figures with ranges. The minimum-viable checklist:

- Could an outsider re-identify the person from your prompt? If yes, redact more.
- Are you pasting an entire document? Summarize first.

4. Logging and audits. Our [Vendor] account logs all prompts. Logs are reviewed [quarterly / when an incident is suspected]. Don't paste anything you wouldn't want reviewed.

5. Escalation. If you're not sure whether something is OK to paste, ask [Owner / your manager] BEFORE pasting. There's no penalty for asking. There's a real penalty for guessing wrong.

6. Reporting an incident. If you accidentally pasted something you shouldn't have, tell [Owner] immediately. Speed of disclosure matters more than the original mistake.

Last updated: [Date]. **Next review:** [Date].

That's it. One page. Maybe 250 words. Answers the 80% of real questions without burying anyone in legal language.

What the policy does and doesn't do

Does:

- Set clear expectations.
- Make the four-class distinction operational.
- Give people permission to ask before pasting.
- Establish reporting / escalation paths.

Doesn't:

- Replace SOC 2 audit work.
- Replace contract review with vendors.
- Cover every edge case (which is fine -- edge cases go to escalation).

The policy is the floor, not the ceiling. Most companies need this floor and don't have it.

SECTION 7

What goes in a SOC 2 conversation

For operators who get audited or whose customers ask about AI compliance posture.

If your business is in a regulated industry, or if you sell to enterprise customers, you'll eventually have a conversation with an auditor or a procurement team about your AI usage. Five questions you should be ready to answer:

> **The five questions auditors and procurement teams ask about AI:**

1. **Which AI tools are in scope?** *Answer: a*
   internal work, [Vendor] Enterprise for client-facing work."
2. **What's the data flow?** *Answer: a*
   team prompts use [Vendor Team]. Client-data prompts route through [Vendor Enterprise]. No customer data flows through consumer-tier accounts."
3. **Where's the DPA?** *Answer: c*
   vendor. Have a copy ready.
4. **Who has access?** *Answer: th*
   access-control system, not maintained separately.
5. **Where's the policy?** *Answer: th*
   by the policy owner.

If you can answer these five with documents in hand, the SOC 2 / procurement conversation is short. If you can't, it's long.

You don't need a 50-page policy document. You need the right 1 page and the right contractual relationship with your AI vendors. Most small operators massively over-engineer this and then still miss the basics.

SECTION 8

Where to go from here

You've now finished Tier 2. Six modules:

7. **Strap In: Train Your Team On AI** *(shipped e*
8. **The Local Comp Kit** *(shipped e*
9. **Role-specific prompt templates** -- building the durable library underneath the rollout.
10. **Multi-step task chains** -- when one prompt isn't enough.
11. **Picking the right model** -- cost-per-task economics from a real operator.
12. **Privacy and what not to paste** *(this one)*
    lets the rest of the curriculum hold up.

You now have the foundation (Tier 1 - Personal) and the operating procedure for using AI at work (Tier 2 - Professional). The next tier -- Tier 3 (Employable) -- is for operators ready to *lead* AI adoption inside an organization. Free primers per module, paid workshops in the cohort. Two primers already shipped (*Your AI Doesn't Know Your Business* and *Your AI B Pretty*); four more drafting.

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 me to come install this whole stack at your business -- the prompt library, the chain workflows, the right-sized model mix, the AI usage policy, the SOC 2-ready answers -- that's the consulting offer. I do the same install I run on my own carpentry business. Done with you, not lectured at you. Real running systems, not whiteboard architecture.

Reach out: alexanderjahn79@icloud.com

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

Closing -- the lock-in line

The single sentence that summarizes this whole module:

4

Four classes of workplace data: public, internal, confidential, regulated.

Each one constrains which AI tier you can use. The 1-page policy makes the distinction operational. The redaction workflow keeps you out of trouble between the lines. Without these, your team will leak something within a year. With them, you compound.

You have the foundation, the operating procedure, and the policy. Tier 3 is where you become the AI-fluent person on a team. See you there.

Agent Logic --
Fond du Lac, WI. This is module 6 of 6 in Tier 2 (Professional).
(Employable) starts next.

theaiguyw
Tier 2 cor