



AI Acceptable Use Policy

Template for [Company Name]

Effective date	[YYYY-MM-DD]
Version	1.0
Owner	[Policy owner — typically General Counsel, COO, or Managing Partner]
Review cadence	Annual, plus on triggered events (new vendor, regulatory change, incident, leadership change)

Before you use this template. Replace every item in square brackets with your own details before employees acknowledge this policy. This template is general information, not legal advice. Have your own counsel review it before you adopt it, especially the disclosure, data handling, and discipline sections.

1. Purpose and scope

This policy governs the use of artificial intelligence (AI) tools by employees, contractors, and other authorized users of [Company Name] (referred to here as "the company") in the course of their work for the company.

It establishes:

- Which AI tools are approved for use, and for what kinds of work
- Which data classes may be used with which tools
- The expectations for human review of AI-produced output
- The consequences of policy violations
- The reporting and incident-response paths

The policy applies to all use of AI tools, whether on company-owned devices, personal devices, on the company network, or elsewhere, when the use relates to company work or involves company-handled information.

2. Definitions

Artificial Intelligence (AI) tool: any software service whose primary function uses large language models, generative AI, predictive AI, or similar techniques to produce outputs (text, code, images, audio, video, data analysis) based on inputs. Includes but is not limited to ChatGPT, Microsoft Copilot, Google Gemini, Anthropic Claude, GitHub Copilot, and any vendor-embedded AI features.

Approved AI tool: an AI tool listed on the company's Approved Tool Registry (the "Registry") with its specific approved-use scope.

Personal AI account: an AI tool subscription, login, or free-tier access controlled by an individual rather than by the company. Personal accounts are not "approved" regardless of whether the underlying tool appears on the Registry.

Data classification: the company's four-tier classification — Public, Internal, Confidential, Restricted. The Data Classification and Handling Guide (separate document) details each class.

Human review: explicit examination of an AI tool's output by a named, accountable employee before the output is used to make a decision, take an action, or leave the company's control.

Shadow AI: any use of AI tools outside the Approved Tool Registry, including use of personal AI accounts for company work.

3. Approved tools

The Registry lists every AI tool approved for company use, the data classes each tool may process, and the specific roles or departments authorized to use it.

- The Registry is maintained by [Governance role — typically the IT director, COO, or designated AI lead].
- The Registry is reviewed quarterly and updated when new tools are evaluated or existing tools change material features.
- A tool not on the Registry is not approved, even if it's a well-known product.

Employees may request additions to the Registry via [how employees request a new tool — for example, a ticket or an email to the governance role]. Requests are evaluated against a vendor due diligence checklist before approval.

4. Approved use cases

[Insert the categories of work AI tools may be used for at your company. The examples below are a starting point.]

- Drafting first-pass internal documents, with human review before any external sharing
- Summarizing documents the employee already has authorized access to
- Researching publicly-available information
- Drafting non-binding client communications (with human review)
- Code generation for non-production systems
- [Add or remove categories to match your company's actual practices]

Use cases not on this list require approval from [Governance role] before AI tools are applied.

5. Prohibited uses

Regardless of which AI tool is involved, the following uses are prohibited:

- **Use of personal AI accounts for any company work.** Free-tier ChatGPT, personal Claude.ai, personal Gemini, personal Copilot Pro, and any other AI subscription not under the company's control may not be used for company work — even on personal devices, even on personal time, when the work involves company information.

- **Submitting Confidential or Restricted data** to any AI tool except those explicitly approved on the Registry for that data class.
- **Allowing AI output to leave the company** without human review when the output represents company work product (client deliverables, regulatory filings, advice, binding communications).
- **Generating output that misrepresents AI involvement** — see Section 8 (Disclosure).
- **Bypassing the company's IT controls** to use AI tools (for example, a personal hotspot to evade network monitoring, or a personal device for company AI work).
- **Using AI tools to make hiring, firing, promotion, compensation, or disciplinary decisions** without human deliberation. AI may inform; it does not decide.
- [Add industry-specific prohibitions here. Example: "no AI on attorney-client privileged communications without partner approval and use of the company's Microsoft Copilot tenant"]

6. Data handling rules by classification

The Data Classification and Handling Guide gives the detail. Summary table:

Data class	Personal AI accounts	Public AI services (free)	Approved tools, default config	Approved tools, restricted config
Public	Permitted	Permitted	Permitted	Permitted
Internal	Prohibited	Prohibited	Permitted	Permitted
Confidential	Prohibited	Prohibited	Permitted only if approved for Confidential	Permitted
Restricted (PHI / ITAR / privileged / regulated)	Prohibited	Prohibited	Prohibited	Permitted only if approved for Restricted

The Registry indicates which classification each tool is approved for.

7. Human-in-the-loop requirements

AI tools produce drafts and suggestions. Humans are responsible for the work product the company delivers.

For any AI output that will:

- Be sent to a client, regulator, court, customer, vendor, or other external party
- Be filed in a system of record (project or case file, customer record, financial system)
- Be used in a billable deliverable
- Affect a personnel decision
- Affect a financial decision above [dollar threshold]

a named, accountable human reviews the output before it is used.

Review means more than "I looked at it." Review means the reviewer is accountable for the content and has verified the facts, the citations (if any), the tone, and the conclusions.

8. Disclosure requirements

[This section depends on your industry. The examples below are a starting point; adopt or modify them based on your regulatory environment.]

The company discloses AI involvement when:

- A regulatory requirement requires disclosure (for example, state bar AI opinions, AICPA ethics rules, FDA, or FTC AI rules where applicable)
- A client's contract requires it
- A client asks
- The output is sufficiently AI-determinative that a reasonable person reviewing the work product would benefit from knowing AI was involved

Disclosure language is provided in the company's standard engagement letter addendum or client communication templates. Employees do not draft ad-hoc disclosures; they use approved templates.

9. Training requirements

All employees complete the company's AI training:

- **New hires:** during onboarding, before any AI tool access is granted (30-minute module and policy acknowledgment)
- **Existing employees:** annual refresher (30 minutes)
- **Triggered:** on any material policy update, all employees complete a delta-training within 60 days

The training covers: this policy, the data classification system, the Approved Tool Registry, common pitfalls (hallucination, confidentiality breach, over-reliance), reporting paths, and case examples (sanitized).

Completion is tracked by [Training owner — typically HR or Compliance].

10. Reporting violations and near-misses

Employees report:

- **Suspected violations** (their own or others') — via [Reporting path — typically an HR portal, an anonymous hotline, or a designated officer]
- **Near-misses** (situations where a violation almost happened) — same path
- **Suspected data exposure** through an AI tool — directly to [Security lead] immediately, in addition to the standard reporting

Reporting in good faith is protected; retaliation is itself a violation of this policy.

Suspected violations are reviewed under the company's Incident Response Policy (separate document).

11. Discipline

Violations may result in discipline ranging from re-training to termination, depending on severity, intent, and recurrence. The company's Disciplinary Procedure (separate document) governs.

In particular:

- **First-time, accidental, no data exposure:** typically retraining and documented acknowledgment
- **First-time, with data exposure:** may include account restriction, disclosure obligations, and disciplinary review
- **Repeat violations:** escalating discipline up to termination
- **Willful violations involving Restricted data:** may include termination, regulatory disclosure, and legal action

12. Review cadence

This policy is reviewed at minimum once per year by [Governance role] with input from [outside or internal counsel]. Triggered reviews occur on:

- A new AI tool added to or removed from the Registry that affects approved use cases
- A regulatory change relevant to the company's AI use (state bar opinion, AICPA ethics ruling, sector-specific rule)
- An AI-related incident at the company
- An AI-related incident at a peer company that the company decides to learn from
- A leadership change at the company

Material updates require employee re-acknowledgment within 60 days.

13. Acknowledgment

I have read and understood this policy. I agree to comply with it. I understand that violations may result in disciplinary action up to and including termination.

I understand that the policy may be updated, and I agree to comply with updated versions on reasonable notice.

Employee name

Signature

Date

Appendices

These appendices are referenced throughout the policy. Keep them with the policy and update them when your tools or data rules change.

Appendix A — Quick decision aid

A one-page decision aid employees can check before using an AI tool. Build it from Sections 3 through 7 of this policy. At a minimum it should ask:

1. Is the tool on the Approved Tool Registry, and am I using a company-controlled account? If not, stop.
2. What class of data am I about to enter, and is this tool approved for that class?
3. Is this use case approved in Section 4? If not, ask [Governance role] first.
4. Will the output leave the company or inform a decision? If so, a named person must review it first.

When in doubt, ask [Governance role] before you proceed.

Appendix B — Approved Tool Registry

The current Approved Tool Registry. It lives at [Registry location — for example, a SharePoint or shared drive link].

Appendix C — Data Classification and Handling Guide

The four-tier classification (Public, Internal, Confidential, Restricted) with examples. It lives at [Data classification guide location].