

RESOURCE · FREE TEMPLATE

AI Acceptable Use Policy: Starter Template

Version v1.0 · MaxtDesign AI Studios

This is a starter acceptable-use policy for generative AI tools at work. It is intentionally concise. The goal is a document every employee will actually read, not one that covers every edge case. Where you see [square-bracket placeholders] swap in your organisation's specifics. Where a section does not apply to your industry, delete it. This template is free to adopt, modify, and share. For regulated industries (healthcare, finance, legal), a one-hour customisation session with MaxtDesign AI Studios adds the industry-specific risk language we can't include generically. Reach out at hello@maxtdesign.com.

1. Scope and purpose

Who this applies to and what it is trying to protect.

This policy applies to every [employee, contractor, intern] of [Company Name] when they use any generative-AI tool (public, internal, or embedded in another product) in the course of their work.

Its purpose is to let you use AI confidently where it helps the business, and to protect [Company Name], its customers, and its counterparties from the three failure modes most associated with AI mistakes: confidential data leaking into a tool that does not respect that trust boundary; AI output being treated as fact without verification; and AI use being inconsistent or undocumented across the team.

2. Definitions

The terms we use throughout the policy.

- **AI tool.** Any software that uses machine-learning models to generate text, code, images, audio, or decisions in response to a prompt. Includes ChatGPT, Claude, Gemini, Copilot, Midjourney, Notion AI, and any feature labelled "AI" inside a tool you already use.
- **Public AI tool.** An AI tool available without an enterprise agreement that binds the provider to data-handling commitments (no-training, no-retention, audit rights). Consumer-tier ChatGPT, consumer Gemini, and most free AI tools fall here.
- **Enterprise AI tool.** An AI tool provided to [Company Name] under a contract that explicitly covers data handling, typically with a no-training-on-inputs commitment and a defined data-retention period. The approved list is in Appendix A.

- **Confidential information.** Anything [Company Name] would not share on its public marketing site. See §4 for the specific data categories.
- **Shipping.** Sending, publishing, or committing AI-generated work in a way that affects a customer, counterparty, or the public. Drafts you keep to yourself are not shipping.

3. Acceptable use

Green, yellow, and red categories. Use this as your day-to-day reference.

We encourage AI use for the kinds of work it is genuinely good at. We require the extra steps in §5 and §6 for anything that reaches customers or touches regulated data.

- 01 Green. Encouraged.** Drafting internal documents, brainstorming, summarising non-confidential reading, learning a new tool or domain, generating unit-test scaffolds from non-confidential code, translating between natural languages, code review suggestions on non-confidential repositories.
- 02 Yellow. Allowed with the §5 verification step.** Drafting customer-facing emails, writing public marketing copy, generating data analysis from anonymised datasets, producing first-pass code that will ship to a non-regulated product.
- 03 Red. Not allowed in any public AI tool.** Pasting or describing in detail: customer PII, unreleased product or financial information, counterparty contracts, credentials, security incident details, legal advice requests, anything covered by a signed NDA, anything under active HR or legal review. Any of these may be processed in an approved enterprise AI tool (Appendix A) if the same rules on verification apply.

4. Data classification: which categories go where

The single most important section. Adjust the categories to match your org's existing data taxonomy.

If your organisation already has a data-classification scheme (Public, Internal, Confidential, Restricted is common), use it. Otherwise, start with the four tiers below.

- 01 Public.** Information [Company Name] has already published: marketing copy, docs, case studies, patents, job listings. Safe in any AI tool.
- 02 Internal.** Information you would share in a company-wide all-hands but not publicly: internal processes, org charts, non-sensitive project updates, internal meeting notes. Safe in enterprise AI tools (Appendix A). Avoid public AI tools unless the content is also de-identified.
- 03 Confidential.** Information shared under a commercial trust expectation: customer data, counterparty contracts, unreleased product plans, unpublished financials, strategy memos. Enterprise AI tools only. Public AI tools are not acceptable, even if redacted, because redaction quality is inconsistent under time pressure.

- 04 **Restricted.** Information covered by a regulatory, legal, or contractual obligation: PHI, payment-card data, material non-public information, active legal-hold material, anything flagged by [Legal, Compliance, Security]. No AI tool, public or enterprise, unless [Security or Compliance] has explicitly approved the specific use case in writing.

GUIDANCE

When in doubt, treat the data as one tier more sensitive than you think. The cost of being cautious is ten minutes of friction. The cost of being wrong is a disclosure letter to affected customers.

5. Verification before shipping

Every AI output that leaves [Company Name] has a human review step.

AI output looks confident whether it is right or wrong. This is the single most-reported failure mode in post-incident reviews and the easiest to prevent.

- **Specific facts.** Numbers, dates, names, citations, quotes attributed to a real person must be verified against a primary source before the output is shipped. "The AI said so" is not verification.
- **Summaries of source documents** must be spot-checked against the source. If the summary is more than two paragraphs, spot-check at least three claims.
- **Code.** If AI-generated code will run in production, it must pass the same review bar as human-authored code: tests, code review, static analysis. AI authorship is not a shortcut around that bar.
- **Charts and visualisations** generated from your data must be cross-checked against the source data on at least three randomly selected data points.

6. Attribution and disclosure

When AI-generated content must be labelled as such.

- External marketing copy, thought-leadership writing, and op-eds published under a named [Company Name] author must disclose AI assistance if the AI did more than grammar or tone editing.
- Customer-facing communications do not require AI disclosure. Employees are professionally responsible for their output regardless of how it was drafted.
- Code committed to [Company Name] repositories should be attributable to the human committer, following the team's existing review process. Internal convention on whether to note AI involvement in commit messages is left to each team.

- Any claim that a [Company Name] product or feature uses AI externally must be reviewed by [Marketing + Legal] before publication.

7. Governance

Who owns this policy, how it updates, and where to raise concerns.

- **Policy owner.** [Head of Security or named person]. Reviews and updates this policy at least annually, or sooner when a tool, regulation, or incident warrants it.
- **Approved tools list** (Appendix A). Updated by the policy owner. Teams may propose additions by emailing [security@company.com].
- **Employee questions** about whether a specific use case is acceptable go to [the policy owner or a designated AI-governance channel].
- **Questions that need legal or privacy review** go to [General Counsel or Privacy Officer]. In regulated industries, these should be time-boxed (e.g. 48-hour response target).

8. Incidents and exceptions

Blameless reporting protects the organisation. Silent fixes guarantee the next incident.

If you realise you have shipped something with an AI-introduced error (a wrong number, a fabricated citation, a misattributed quote), or if you have pasted something into a public AI tool that fell into the Confidential or Restricted categories, report it to [the policy owner] immediately.

Reporting is blameless. The purpose is to assess impact, fix the output if it reached a customer, and identify whether the policy or training needs an update. Retaliation against a person who self-reports in good faith is itself a policy violation.

- One-off exceptions to §3 (Acceptable Use) or §4 (Data Classification) can be granted by [the policy owner or compliance lead] in writing, case-by-case, with a dated expiry.
- Standing exceptions are not granted. If a use case needs a permanent carve-out, the policy itself should be updated.

9. Training

Everyone using AI for work completes a short onboarding.

All new [employees, contractors] complete an AI acceptable-use training as part of onboarding. The training covers this policy plus the high-value skills that prevent the most common mistakes: prompt structuring, verification habits, and data-classification judgment under time pressure.

Existing [employees, contractors] complete a refresher annually or when this policy materially changes. Teams with elevated exposure (customer support, legal, finance, engineering touching customer data) may have additional role-specific training.

Appendix A. Approved enterprise AI tools

Fill in with your organisation's enterprise-licensed AI stack. Update when contracts change.

The following AI tools have been reviewed for data handling and are approved for use with data up to and including the Confidential tier. Restricted-tier data requires additional per-use approval from [Security or Compliance] regardless of tool.

- **[ChatGPT Enterprise / Claude for Enterprise / Gemini for Workspace]**. Approved for general text drafting, summarisation, Q&A up to Confidential. [Start date] through [end of contract]. Data-handling commitment: [no training on inputs, N-day retention].
- **[GitHub Copilot Enterprise / Cursor / etc.]**. Approved for coding assistance on [Company Name] repositories up to Confidential. Data-handling: [as above].
- **[Add your tools here]**

Appendix B. Change log

Every material change to this policy is recorded here with date, author, and a one-sentence summary of what changed and why. The living history matters more for audit purposes than a clean version number.

- [YYYY-MM-DD]. [Policy owner]. Initial adoption of MaxtDesign AI Studios starter template.