Gen AI Acceptable Use Policy - Draft

1. Purpose

The purpose of this policy is to establish guidelines for the responsible and secure use of Generative Artificial Intelligence (GenAI) tools (e.g., OpenAI ChatGPT, Microsoft Copilot, Anthropic Claude, Midjourney) within **[Company Name]**. This policy aims to leverage AI innovation while ensuring the protection of intellectual property, client confidentiality, and compliance with regulatory frameworks including **CMMC, NIST SP 800-171, and HIPAA**.

2. Scope

This policy applies to all employees, contractors, consultants, and third-party partners performing work on behalf of **[Company Name]**. It covers all GenAI tools accessed via web interfaces, API integrations, browser extensions, or embedded software features.

3. General Usage Guidelines

- **Authorized Tools Only:** Employees may only use GenAI tools that have been explicitly approved by the Security Team and listed in the **Authorized AI Service Catalog**.
- **Human Oversight:** AI-generated content (code, text, or data analysis) must not be treated as a final product. All output must be reviewed, validated, and "owned" by a human subject matter expert before being used for business decisions or client deliverables.
- **Transparency:** Any significant work product generated or heavily assisted by GenAI must be disclosed to the internal supervisor or, where contractually required, to the client.

4. Protection of Sensitive Information (The "Zero Trust" Rule)

The "context window" of a GenAI tool is effectively a flat namespace with no internal access control. Information shared with a public GenAI model must be considered public information.

- **Public/Consumer AI Prohibitions:** Under no circumstances shall the following be entered into consumer-grade or "free" AI tools:
  - **Intellectual Property:** Proprietary algorithms, product roadmaps, or unreleased research.
  - **Internal Credentials:** API keys, passwords, or system configurations.
  - **Client Data:** Any data belonging to a client that is not already in the public domain.
- **Data Scrubbing:** Before using authorized AI tools, users must ensure all input data is scrubbed of sensitive identifiers.

5. Regulatory Compliance Addenda

A. CMMC & NIST SP 800-171 (Defense Industrial Base)

For projects involving Federal Contract Information (FCI) or Controlled Unclassified Information (CUI):

1. **Authorized Boundaries:** CUI and FCI must only be processed within GenAI tools that reside within a FedRAMP Moderate (or higher) authorized boundary (e.g., Azure OpenAI in Azure Government).
2. **Prohibition of Consumer AI:** Use of standard ChatGPT, Claude, or Gemini for any contract-related work is strictly prohibited.
3. **Data Lineage:** Any AI-assisted processing of CUI must be documented in the System Security Plan (SSP) to ensure auditability during CMMC assessments.
4. **No Training on CUI:** AI service contracts must explicitly state that input data (prompts) will not be used to train the provider's global models.

B. HIPAA (Healthcare & Life Sciences)

For projects involving Protected Health Information (PHI):

1. **Business Associate Agreements (BAA):** No PHI may be entered into an AI tool unless **[Company Name]** has a signed BAA in place with the AI service provider.
2. **Minimum Necessary Standard:** AI tools must only be granted access to the minimum PHI required to perform the intended task.
3. **De-identification:** Preference must be given to using de-identified data or synthetic datasets for AI-driven analysis.
4. **Right to Erasure:** AI service providers must demonstrate the technical capability to fulfill "Right to be Forgotten" requests by removing specific data from their logs and vector stores.

6. API Integrations & Agentic AI

- **Security Review:** All requests to integrate AI APIs into internal applications must undergo a formal security impact analysis.
- **Least Privilege:** AI Agents (autonomous bots) must be assigned unique identities and restricted permissions. Agents must never have "standing privileges" to databases; they must use Just-in-Time (JIT) access tokens.
- **Sandboxing:** AI-generated code must be tested in a dedicated, isolated sandbox environment before being promoted to production.

7. Shadow AI & Browser Extensions

- **Extension Blocklist:** AI-powered browser extensions that require "read all site data" permissions are prohibited on managed workstations unless approved by the CISO.
- **Monitoring: [Company Name]** reserves the right to monitor network traffic for

unsanctioned AI domain access to identify and remediate "Shadow AI" usage.

8. Incident Reporting

Any suspicion that sensitive data has been accidentally shared with an unauthorized AI tool must be reported to the Security Team immediately as a **Data Spillage Incident**.

9. Enforcement

Failure to comply with this policy may result in the revocation of AI access, disciplinary action, and, in cases involving regulated data (CUI/PHI), mandatory reporting to federal authorities.

Policy Version: 1.0

Effective Date: March 22, 2026

Approved By: [Name/Title]