Blue Sheen

AI acceptable use policy

Aligned to NIST AI RMF, ISO/IEC 42001, and the EU AI Act

May 23, 2026

Prepared for Acme Mid-Market Co.

Executive summary

This policy applies to all 312 employees and contractors of Acme Mid-Market Co. (“Acme”) who use, evaluate, or build AI systems on behalf of the company. It is aligned to NIST AI RMF 1.0, NIST AI 600-1 Generative AI Profile, ISO/IEC 42001:2023, and the EU AI Act (Regulation 2024/1689) for orgs with EU customers in their roadmap.

Acme operates a HIPAA-covered SaaS platform serving healthcare clinics in the United States, holds SOC 2 Type II certification, and is in active GDPR scoping. The policy is calibrated to those realities.

The policy explicitly addresses three vectors that mid-market healthcare SaaS companies most frequently misjudge: PHI exposure through prompts, shadow AI in clinical contexts, and vendor sub-processor risk from AI-augmented tools that previously didn’t process PHI.

This policy is not a substitute for legal counsel review. Sections flagged with “[GC review]” require sign-off from the Acme General Counsel before adoption.

Scope and applicability

This policy covers:

- All employees, contractors, and interns of Acme who access AI systems for work
- Third-party vendors whose contracts require AI usage compliance
- Internal AI systems Acme builds or deploys
- External AI services Acme uses (ChatGPT Team, GitHub Copilot Enterprise, and any future additions)

It does NOT cover:

- Personal AI use on personal devices outside Acme work
- AI systems used by customers within their own Acme tenant (covered by customer terms)
- Academic or research uses already governed by separate IRB protocols

Defined terms

AI system: a machine-based system that, for explicit or implicit objectives, infers from the input it receives how to generate outputs such as predictions, content, recommendations, or

decisions that can influence physical or virtual environments. (Adapted from EU AI Act Article 3.)

Generative AI: an AI system specifically intended to generate, with varying levels of autonomy, content such as text, images, audio, video, or code.

PHI: Protected Health Information as defined by HIPAA — individually identifiable health information held or transmitted by a covered entity or business associate.

High-stakes AI use: any AI-supported decision or output that (a) directly affects a patient's clinical care, (b) is shared with a customer without human review, or (c) is used as evidence in a contractual, legal, or regulatory matter.

Approved AI tools

Tool	Approved for	Restrictions
ChatGPT Team (60 seats)	General knowledge work, drafting, brainstorming	NEVER paste PHI. Use the team workspace, not personal accounts.
GitHub Copilot Enterprise (45 dev seats)	Code completion, code review assistance	Code involving auth, billing, or PHI handling must be human-reviewed before merge.
Claude (any plan)	Currently NOT approved	Pending procurement review (Q3 2026 target).

Any new AI tool must be reviewed by the AI Governance Committee (see Roles & Responsibilities) before use. Personal AI accounts on personal devices for Acme work are prohibited.

Data classification and handling

Acme uses four data tiers. AI handling rules per tier:

Tier 1: Public. Marketing copy, published help articles, public-facing docs. May be freely shared with approved AI tools.

Tier 2: Internal. Internal docs, project plans, non-customer-specific analysis. May be shared with approved AI tools in the team workspace only.

Tier 3: Confidential. Customer-identifiable info that is NOT PHI (e.g., customer org name + non-clinical usage data). Allowed only with the customer's BAA-compatible AI tool list verified. Default: do not use AI on Tier 3 without AI Governance Committee approval.

Tier 4: Restricted (PHI / financial / IP). NEVER paste, upload, or reference in any AI tool — including approved ones — unless the tool is explicitly under a BAA with Acme. Today no tool meets that bar. Until further notice, Tier 4 = AI-prohibited.

Acceptable use

Approved AI use cases:

- Drafting marketing and internal communications (Tier 1 and 2 inputs only)
- Code completion and code review on non-Tier 4 codebases
- Research, summarization, and idea generation
- Meeting notes and prep (no PHI mentioned)
- Customer support draft responses for human review (no PHI in prompt)

Prohibited use

The following are explicitly prohibited:

- AI use in clinical decision support shown to customers without explicit clinical review
- Pasting PHI into any AI tool, including approved ones, regardless of tier
- Using AI to generate evaluations, performance reviews, or compensation decisions about individual employees without human approval
- Auto-sending AI-generated email or chat responses to customers without human review
- Training external AI models on Acme data, customer data, or PHI

Transparency and disclosure

Acme discloses AI assistance for high-stakes outputs only. This includes:

- Any customer-facing report that was substantially AI-drafted
- Any clinical-context content reviewed by AI before human approval
- Marketing copy at a customer's explicit request

Routine internal drafts, code completions, and brainstorming do not require disclosure.

Incident reporting

Any of the following triggers a same-day incident report to the AI Governance Committee:

- PHI accidentally pasted into any AI tool
- AI-generated output sent to a customer without human review
- A customer asking whether AI is used in our product (this is a notable signal)
- A new AI tool installed without committee approval (shadow AI)
- Any AI output that contradicts a clinical guideline or a known regulatory standard

Acme's existing incident response process (documented but not yet tested per the wizard intake) applies to AI incidents until a dedicated AI incident playbook is built (Q3 2026).

Roles and responsibilities

Role	Responsibility
CEO	Owns the policy. Approves exceptions in writing.
CTO + General Counsel	Co-chair the AI Governance Committee (NEW — to be formed within 30 days of policy adoption).
AI Governance Committee	Reviews new tool requests, owns the approved-tools list, reviews incidents, owns the annual policy refresh.
Department managers	Ensure their teams complete annual training. Surface shadow AI to the committee.
Every employee	Compliance with this policy. Reporting incidents.

Training

All employees complete a 20-minute AI policy training within 30 days of hire and annually thereafter. Engineering and Customer Success roles complete an additional 30-minute role-specific module covering Copilot guardrails (eng) and customer-disclosure rules (CS).

Vendor risk management

Any vendor whose product or contract adds AI capabilities must be reviewed by the AI Governance Committee before Acme adopts the new feature. Specifically:

- A previously-non-AI vendor announcing AI features → Committee review before enabling
- A new vendor whose product is AI-first → Committee review pre-contract
- BAA review extension for any AI vendor touching PHI

Audit, review, and exceptions

This policy is reviewed at least annually by the AI Governance Committee and after any material AI incident. Exceptions require written CEO approval. Reviews are logged and retained for 7 years per Acme's broader compliance retention policy.

References and alignment

This policy aligns to the following published frameworks. Section-by-section mapping is maintained in the AI Governance Committee shared workspace.

- **NIST AI RMF 1.0** (January 2023): Govern, Map, Measure, Manage functions
- **NIST AI 600-1** (July 2024): Generative AI Profile, especially actions GV-1.1, MP-2.3, MS-2.5, MG-3.2
- **ISO/IEC 42001:2023**: AI management system requirements (target: full conformance by EOY 2027)
- **EU AI Act (Regulation 2024/1689)**: General-purpose AI obligations (Article 51-55), even though Acme has no current EU customers — preparatory
- **HIPAA Security Rule**: 45 CFR 164.308 and 164.312 administrative + technical safeguards
- **SOC 2 Trust Services Criteria**: especially CC6 (Logical Access) and CC7 (System Operations)

Appendix A: Rollout checklist

- Week 1: CEO sign-off; AI Governance Committee formed (CTO + GC co-chair)
- Week 2: All-hands email + 20-min training module published
- Week 3: Department managers brief their teams; shadow-AI amnesty window opens

- Week 4: Tools list audit; non-approved AI tools removed or flagged
- Week 5: First quarterly review scheduled; incident-reporting workflow tested
- Week 6: Policy linked in onboarding docs; customer-disclosure update reviewed by GC

Appendix B: Acknowledgment

By signing below, I acknowledge I have read, understood, and agree to comply with the Acme Mid-Market Co. AI Acceptable Use Policy.

Name: _____ Role: _____ Date: _____ Signature: _____

This sample policy was hand-crafted by Blue Sheen for a fictional client. Your real policy will be tailored to your jurisdiction, frameworks, AI maturity, and risk appetite. Request your custom policy at bluesheen.com/tools/ai-policy-generator/.