GRAYPOINT TECHNOLOGY

AI Vision & Ethics Policy Template

A starter framework executive teams can use to define their AI vision, set ethical guardrails, and build the governance policy that puts it into practice.

> **How to use this template**
>
> This template is a starting point, not a finished policy. Work through it section by section with your leadership team, replace every bracketed prompt with your own language, and delete these instructions before distributing the final document.
>
> It is written for small and mid-sized organizations adopting AI tools internally or client-facing. *It does not constitute legal advice — have qualified counsel review your final policy against the laws and regulations that apply to your industry.*

1. What Is an AI Vision — and Why It Comes First

An AI vision is a short, plain-language statement of why your organization uses AI, what value it is meant to create, and the standards it will always be held to. It is the "why" that every later rule, guardrail, and approval process traces back to.

Without a vision, ethics and policy work becomes a list of restrictions with no shared purpose behind them — hard to explain to staff, harder to apply consistently to a new AI tool nobody anticipated. With a vision in place, every later policy decision becomes a test of fit: *"Does this use serve the vision, and does it respect the standards we committed to?"*

Your AI Vision Statement

Draft your vision (2–4 sentences):

> *Example starting point: "[Organization] uses AI to help our people work faster and serve clients better — never to replace human judgment on decisions that affect people's rights, safety, finances, or trust in us. Every AI-assisted output is reviewed by a qualified person before it reaches a client, employee, or the public."*

2. Guiding Ethical Principles

Select and adapt the principles that matter most to your organization. Most policies rest on some version of the six below.

- Human oversight & accountability — a qualified person reviews and remains responsible for AI-assisted outputs before they are acted on or shared externally.
- Fairness & non-discrimination — AI use is monitored for bias that could disadvantage employees, clients, or applicants.
- Transparency & explainability — people are told when AI materially shaped a decision or deliverable that affects them, and can ask how.
- Privacy & data minimization — only the data necessary for the task is collected, used, or retained by an AI system.
- Security by design — AI tools and the data they touch are protected with encryption, access controls, and vendor due diligence.
- Reliability & safety — AI outputs are validated before use in decisions with legal, financial, safety, or reputational consequences.

Principles we are adding or adapting for our organization:

> *[List any industry-specific principles — e.g., HIPAA, GLBA, or client-confidentiality commitments unique to your business.]*

3. Scope & Definitions

This policy applies to (check all that apply):

- ☐ Generative AI writing/research tools (e.g., ChatGPT, Claude, Gemini, Copilot Chat)
- ☐ AI features embedded in existing software (CRM, help desk, Microsoft 365 Copilot, etc.)
- ☐ AI-based automation or autonomous agents
- ☐ AI used in hiring, performance, or personnel decisions
- ☐ Client-facing AI features or chatbots
- ☐ AI used in security, fraud detection, or risk scoring

Other (list any additional tools, platforms, or use cases not covered above):

> *[Add anything specific to your organization that is not already checked off above.]*

This policy does NOT cover (check all that apply):

- ☐ Basic spam filtering or spell-check with no material decision impact
- ☐ Built-in productivity features (e.g., autocomplete, grammar suggestions) with no material decision impact
- ☐ AI features disclosed and excluded under a separate policy (e.g., a dedicated hiring-AI or security-AI policy)

Other exclusions (list and explain why they are out of scope):

> *[e.g., basic spam filtering, spell-check, or other AI features with no material decision impact — list your exclusions and why.]*

4. Permitted & Prohibited Uses

Use this table as a working draft — most organizations find it easier to start from prohibited uses, then define what is left as permitted.

Permitted Uses	Prohibited Uses
[e.g., drafting internal documents for human review]	*[e.g., final hiring, credit, or termination decisions made by AI alone]*
[e.g., summarizing public research or filings]	*[e.g., entering client PII/PHI into non-approved public AI tools]*
[Add your own row]	*[Add your own row]*

5. Data Governance & Privacy Commitments

- Data minimization — collect and submit only what is necessary for the specific AI task.
- Encryption — data is encrypted in transit (TLS 1.2+) and at rest wherever it is processed or stored.
- Retention limits — AI inputs/outputs are retained for a defined period, then deleted or reviewed (commonly 30–90 days for working data).
- Data subject rights — a process exists for individuals to ask what data was used and request correction or deletion where applicable.

Our retention period for AI-related data is:

> *[e.g., 30 days for transient working data, unless preserved in an approved business record.]*

Approved AI tools/vendors for handling sensitive data:

> *[List the specific tools your organization has vetted and approved — everything else is prohibited by default.]*

6. Human Oversight & Accountability

Define who is accountable for reviewing AI-assisted work before it leaves the building — this is the single most important control in most AI policies.

Our review requirement:

> *[e.g., "No AI-generated output is sent to a client, published publicly, or used in a personnel/financial decision without review and sign-off from a qualified employee."]*

Escalation path if something looks wrong:

> *[Name the role/person staff should notify, and what happens next.]*

7. Third-Party & Vendor AI Oversight

Apply this checklist before adopting any new AI vendor, platform, or embedded AI feature.

- Vendor security and privacy practices reviewed (SOC 2, security questionnaire, or equivalent).
- Data processing terms reviewed — what the vendor can retain, use for model training, or share.
- Incident notification and breach-response terms confirmed in writing.
- Re-assessment scheduled for higher-risk vendors (e.g., annually, or on contract renewal).

Vendor review owner:

> *[Name/role responsible for vendor due diligence before a new AI tool is approved for use.]*

8. Roles & Responsibilities

Role	Responsibility
Executive Sponsor	*[e.g., owns the AI vision, approves policy changes]*
AI/IT Governance Owner	*[e.g., maintains the approved-tool list, coordinates vendor reviews]*
All Employees	*[e.g., follow permitted-use rules, escalate concerns, complete required training]*

9. Review & Update Cadence

This policy will be reviewed:

> *[e.g., "At least annually, and any time we adopt a new AI tool, change vendors, or a relevant law or regulation changes."]*

10. Approval & Sign-Off

Name / Title	*Date*

Want help turning this into a finished, board-ready AI governance policy? GrayPoint Technology provides vCIO/vCISO advisory to build and operationalize AI, data, and security policy for executive teams. Schedule a consultation at graypointtech.com.