Implementing Trustworthy Artificial Intelligence:

A Practical Policy Template for Organizations

DELINEATE

8275 Allison Pointe Trail,
Suite 130 Indianapolis, IN 46250

Contact Us
317.730.3991
info@delineateconsulting.com

Introduction

This document provides a sample Artificial Intelligence (AI) Policy designed to help organizations begin their journey toward responsible AI adoption and governance. It is not intended to be a final or comprehensive framework but rather a high-level example that outlines core principles, roles, and responsibilities.

The purpose of this sample policy is to:

- Demonstrate how organizations can establish foundational guardrails for AI use.
- Encourage safe, ethical, and compliant application of AI technologies across business functions.
- Provide a starting point for tailoring policies to a company's specific context, industry, and regulatory obligations.

It is important to note what this document is not. It is not a substitute for company-specific policies, regulatory compliance programs, or detailed operational procedures. AI governance requires supplemental policies, standards, and processes in areas such as data management, risk assessment, vendor oversight, system approval, and workforce training. Each organization must adapt and expand upon this framework to reflect its unique mission, values, and risk profile.

We recognize that organizations are at different stages of AI maturity. For some, this policy may serve as an introduction to responsible AI practices. For others, it may highlight gaps in existing governance and spark deeper conversations about risk management, oversight structures, and technical safeguards.

Our goal in sharing this sample policy is to start the discussion. We can help organizations build upon this foundation by developing tailored AI governance policies, implementation procedures, and oversight mechanisms that align with both business objectives and emerging regulatory requirements.

Other Implementation Considerations

When implementing this template:

- Customize the tone and content to match company culture.
- Add specific examples for your industry and company, especially in the encouraged use cases section.
- Update the approved tools section to those you have actually approved.
- Update data governance, AI governance charter structure, supporting procedures, etc. throughout.
- Consider creating a one pager or infographic for quick reference.
- Create a Claude or ChatGPT Project with this policy as a knowledge base for your employees.
- Link out to any training, list of AI Champions and AI Business Partners and the Office Hours schedule

Further adapt based on your:

- Company size and structure
- Geographic location and regulations
- Industry requirements
- Regulatory environment
- Current AI maturity
- Risk tolerance
- Employee technical literacy

Engage with Delineate

Ready to transform your organization with responsible AI? Connect with us and let's make AI your ally in generating impactful solutions that drive positive change.

Contact Us

317.730.3991
info@delineateconsulting.com

Artificial Intelligence Standard

(Last Updated August 2025)

Our AI Policy in Plain Language

We believe AI tools can make you more effective and help us to be more successful as a company. We encourage the use of AI tools and this policy outlines how to do so responsibly. Here are the key points:

We want to enable you to:

1. Use AI to improve your work quality, speed, and outcomes.
2. Learn how to get the most out of AI tools and be a Responsible AI user.
3. Seek help when you're unsure.

Always:

1. Understand the security and privacy of an AI system to keep all customer and company confidential information private.
2. Lead with your subject matter expertise and double check AI-generated content before sharing.
3. Use company AI tools or request new use cases when needed.
4. Share your successful uses of AI with colleagues.

The remainder of this document provides more detailed guidance. We encourage review the document and interacting with the policy using our AI Standard Policy GPT.

Purpose and Scope

Purpose Overview

This section explains why [Company] has adopted an AI governance policy and the objectives it seeks to accomplish. As an organization that embraces the opportunities of AI, our intent is not to restrict innovation but to enable responsible adoption at scale. Maintain an AI positive tone in these statements.

Example text: This policy establishes a structured framework for how [Company] develops, deploys, and oversees AI-powered systems that streamline operations, improve workflows, and enhance enterprise value. It applies to the entire AI lifecycle, spanning data collection, model design, testing, deployment, day-to-day use, and continuous monitoring. This policy covers all employees, contractors, and partners who interact with AI tools in the course of company work.

The purpose is to ensure that AI systems used to identify inefficiencies and recommend process improvements are done so ethically, lawfully, and transparently, with human oversight embedded in decision-making. By aligning with emerging best practices and regulations, this policy protects both [Company] and its stakeholders. Importantly, recent industry research highlights that achieving real business outcomes with AI requires more than just advanced models. It demands proactive, embedded governance processes, roles, and checks that become part of daily workflows rather than a one-time compliance exercise. This policy embodies that approach, ensuring governance is a catalyst for safe innovation, not a barrier.

Scope Definition
Here we define what this policy covers. Include all employees, contractors and potential partners that interact with company data. Also, broadly specify which types of AI tools are covered.

Example text: This policy applies to all individuals working with or on behalf of [Company], including employees, contractors, consultants, and approved external partners, whenever they use AI in the context of company business or when interacting with company data.

The scope includes all categories of AI systems and tools, such as:

- Large Language Models (LLMs): Examples include conversational assistants, chatbots, text generation, summarization, and translation.
- Generative Media Tools: image, audio, and video generation or editing platforms.
- Code and Productivity Assistants: AI tools that generate, review, or optimize code and documentation.
- Automation and Agentic Systems: task automation bots, workflow optimizers, and multi-agent orchestration platforms.
- Analytical and Predictive Models: forecasting, optimization, anomaly detection, and similar statistical and machine-learning–based analytics.

This policy applies both to AI tools currently in use and to new AI systems introduced in the future, ensuring that emerging technologies are incorporated into governance as they are adopted. Any AI solution, whether built internally, licensed from a vendor, or accessed as a cloud service, falls within scope if it processes company data or influences company decision-making.

Our AI Philosophy

This section sets the tone for the entire policy. For an AI-positive tone, make a clear statement about embracing AI as a core enabler of business success.

Example text: At [Company], we see artificial intelligence as a strategic enabler of efficiency, innovation, and growth. We believe AI is most powerful when it augments human judgment and creativity, empowering people to work smarter, faster, and with greater impact. AI is not a substitute for human expertise but rather a tool that extends our capabilities and unlocks new opportunities.

We encourage employees, teams, and partners to experiment responsibly with AI tools, recognizing that early adoption and continuous learning will position us ahead of competitors. By embedding AI into our daily operations, we aim to drive and measure business value while fostering a culture of curiosity and innovation.

Our ambition is to become a Responsible AI-enabled organization that not only uses AI to improve efficiency and decision-making but does so in a way that is ethical, transparent, and aligned with stakeholder trust. We will empower and reward individuals and teams that integrate AI responsibly into their workflows, reinforcing our commitment to both innovation and accountability.

Governance Structure and Accountability

This section provides the structure of AI policy, establishing roles and responsibilities, actors, and the process for AI use-case approval and usage. This may be nested within a larger governance body as a subcommittee or a stand-alone feature. In either case, a separate Governance Charter should be established that minimally defines roles, accountabilities, processes, and outcomes.

Example text: Effective AI governance must establish clear roles and accountability at all levels within [Company]. Leadership treats AI governance as a strategic priority (not just an IT concern), with executive sponsorship ensuring resources and support. Key elements of the governance structure include:

- **AI Governance Committee:** A cross-functional committee or working group to oversee AI projects and risks. This team includes stakeholders from across [Company], such as IT, operations, sales, research and development, compliance/legal, security, and HR. This interdisciplinary AI governance committee collaborates to set priorities, align AI use with business value, ethics, and ensures accountability. The committee will review proposed AI use cases, approve deployments, and establish success criteria for AI outcomes.

- **Defined Roles and Responsibilities:** Every AI project is assigned clear ownership for governance tasks. Each project has designated Responsible AI Leads/Stewards who coordinate risk assessments and mitigation. Data Stewards (as defined in our data governance framework) are responsible for data quality and policy compliance in model datasets or task automation procedures. During deployment, solutions must implement required controls (audit logging, access permissions, etc.), and business managers using the AI must ensure human oversight on AI-driven decisions.

  Additional Consideration: *While many companies report having AI usage policies, fewer have appointed governance roles or incident response playbooks, leading to a gap between policy and practice.*

- **Executive and Board Oversight:** The Board of Directors (or top executives) will receive regular updates on AI initiatives, risks, and compliance status. The board (or a relevant risk committee) should formally approve this AI governance policy and integrate AI risks into enterprise risk management. [Company] leadership will proactively treat AI as a governance pillar, and not simply a tech trend, enabling an AI positive culture that helps achieve our goal of [Company AI mantra].

  Additional Consideration: *Building AI fluency at the leadership level is essential, for example through quarterly AI briefings or dedicated board training sessions on emerging trends and regulatory developments.*

Definitions

To ensure clarity and consistency across this policy, the following terms are defined. These definitions should be referenced whenever technical, ethical, legal, or regulatory terms appear within the policy.

- **Artificial Intelligence (AI):** The application of machine learning, algorithms, automation, and software to perform tasks, recognize patterns, and generate predictions or recommendations based on data and programmed logic.

- **Artificial Intelligence Committee (AI Committee):** An internal [Company] body responsible for reviewing, approving, and overseeing all uses of AI within [Company].

- **Artificial Intelligence System (AI System):** Any software solution that leverages one or more AI techniques to achieve specific human-defined objectives, producing outputs such as content, forecasts, recommendations, or decisions that influence processes or environments.

- **Closed AI System:** An AI model in which input provided by a user is isolated and used only within that user's environment. Input data is not shared across users, making it more secure and private.

- **Embedded AI Tools:** AI functionality that is built into existing, approved enterprise software applications (e.g., AI features within productivity suites) that do not require separate approval by the AI Committee.

- **Non-Public [Company] Data:** Any data or information that, if disclosed without authorization, could: (i) infringe on individual privacy rights; (ii) violate regulations or laws; (iii) compromise [Company]'s financial position; (iv) harm its reputation; or (v) reduce competitive advantage.

- **Open AI System:** An AI system where user inputs may be combined with data from other users to further train or refine the model. Inputs are not considered private and could be visible to third parties or other users.

- **Personal Information:** Any information that directly or indirectly identifies, relates to, describes, or could reasonably be linked to a specific individual or household.

- **[Company] Representatives:** Includes all [Company] directors, officers, board members, employees, contractors, agents, resellers, distributors, affiliates, and any third parties performing services on behalf of [Company].

Encouraged Uses and Core Principles for Responsible AI

It is common to see AI policies focus on what you can't do with AI. While this is an important aspect, it can drive AI utilization underground, Approaching the policy from an AI-positive tone highlights specific ways employees are encouraged to use AI. Leading with what you can do will positively reinforce desired behaviors rather than drive users into uncontrolled AI Systems. Make it clear this list is not exhaustive but illustrative.

Example text: At [Company], we want to highlight the many positive and productive ways AI can be applied to help employees, teams, and partners work more effectively. This section emphasizes how

AI can augment your work, drive innovation, and accelerate results. This list is not exhaustive but is intended to inspire safe and impactful AI adoption.

To maximize the benefits of AI:

1. Complete [Company] AI Training to learn how to design automations, write effective prompts, and engage with AI systems responsibly.
2. Be clear in your objectives when using AI to ensure outputs align with your intended outcomes.
3. Verify accuracy, particularly when outputs involve sensitive, financial, legal, or customer-facing information.
4. Iterate and improve by treating AI as a collaborator that gets better with feedback. Think about it more as a whiteboarding exercise than a final answer.
5. Provide context using tools like [ChatGPT Projects/Claude Projects] to improve relevance and quality.
6. Document effective prompts and workflows and share your story to help colleagues adopt best practices.
7. Start small with pilot tasks before expanding AI use across larger processes.
8. Keep a record of AI-assisted approaches for audit, reuse, and compliance tracking.

Example AI Applications:

Marketing

- Draft and A/B test campaign copy
- Generate creative visuals
- Analyze customer sentiment across channels
- Forecast campaign performance

R&D

- Accelerate literature reviews
- Simulate scenarios
- Design experiments
- Uncover patterns in large datasets

Legal

- Analyze contracts for risk
- Summarize regulatory updates
- Flag deviations from policy
- Assist with due diligence research

HR & Talent

- Draft job descriptions
- Analyze engagement survey results
- Build personalized learning content
- Support workforce planning scenarios

Sales

- Research prospective clients
- Generate tailored outreach messages
- Summarize meeting notes
- Recommend upsell or cross-sell opportunities

Engineering / IT

- Accelerate code development and review
- Assist with debugging
- Generate technical documentation
- Detect anomalies in system logs

Finance

- Support financial close processes
- Automate reconciliations
- Detect anomalies in transactions
- Assist in M&A analysis and integration planning.

Operations & Supply Chain

- Forecast demand
- Optimize routing and scheduling
- Predict inventory shortages
- Analyze supplier performance trends

Approved Tools
Provide a listing and summary of approved applications to each AI System with links to full operating procedure and training. In the table provide a brief summary of the practical application and the scope of the system.

Example text:

Approved AI System	Business Application	Intended Use (Scope)
Operations Optimization Co-Pilot (Core Platform)	Identifies process bottlenecks and recommends workflow improvements across finance, operations, and sales.	Provide optimization recommendations and playbooks; business owners review/approve changes before execution.
AP Invoice Intelligence	OCR/extraction of invoices; GL code suggestions; duplicate/fraud flags.	Assist AP clerks with data capture and coding; no auto-payment or vendor changes without human approval.
FP&A Forecasting Assistant	Revenue/expense/cash forecasting; scenario planning.	Support planning cycles with scenario analysis; forecasts require FP&A sign-off before publication.
Sales Lead Scoring & Next-Best-Action	Prioritizes leads/opportunities and suggests outreach steps based on history/fit.	Coach reps on who to contact and how; tool does not send external communications without rep approval.
Customer Health & Churn Risk	Flags at-risk accounts and expansion signals from product and ticket data.	Triage and playbooks for CSMs; cannot auto-downgrade or reduce service without manager approval.
Marketing Content Co-Writer	Drafts emails/ads/landing copy using brand library and product facts.	Create draft content only; human/legal review required for claims and regulated topics before release.
HR Talent Match (De-identified)	Matches de-identified resumes to job requirements and surfaces candidates.	Triage/screening aid; final hiring decisions remain human; fairness testing and HR oversight required.
IT Helpdesk Virtual Agent	Answers FAQs, gathers diagnostics, routes tickets, triggers safe automations.	Tier-0/1 support; suggests steps and opens tickets; escalates when uncertain or on policy-sensitive issues.
Engineering Code Assistant	Suggests code, tests, and documentation; inline explanations.	Developer assistance; all code undergoes peer review and security/license scanning before merge.
Contract Review Summarizer	Extracts clauses, deviations, and risk from NDAs/MSAs/SOWs.	First-pass legal review; attorneys make determinations and negotiate terms.
Vendor Risk Triage (AI)	Scores vendor security/privacy posture using DDQ and external signals.	Prioritize due diligence; Procurement/Compliance retain approval authority.
Demand & Inventory Optimizer	Forecasts demand and recommends reorder points and safety stock.	Planner decision support; planners review/approve POs; model cannot auto-issue purchase orders.
Predictive Maintenance	Detects equipment anomalies from sensor/telemetry data.	Maintenance scheduling recommendations; cannot bypass safety interlocks or lock-out/tag-out procedures.
Security Threat Triage	Detects anomalous logins, malware, and phishing; prioritizes alerts.	SOC triage aid; cannot auto-block accounts/endpoints without analyst

Approved AI System	Business Application	Intended Use (Scope)
		review (except pre-approved playbooks).
Enterprise Search & Summarization	Secure chat over internal docs with access controls and citation.	Knowledge retrieval for employees with valid access; no export of restricted/PII outside approved channels.
Data Quality & PII Scanner	Identifies PII and data quality issues in data stores and pipelines.	Alert stewards; no auto-deletion or schema changes; remediation tracked in data governance backlog.
Expense Audit Assistant	Flags potential policy violations or duplicates in expense reports.	Auditor triage; employees notified to correct; no auto-denials without human review.
Regulatory Change Monitor	Tracks new/updated regulations and summarizes potential impacts.	Compliance awareness and task generation; not a substitute for legal advice or counsel review.
Contact Center QA Bot	Transcribes calls, scores QA metrics, surfaces coaching moments.	Supervisor coaching insights; not sole basis for disciplinary action; HR review for material decisions.
Document Classifier & Router	Classifies and routes inbound forms/emails to correct queues.	Work routing only; cannot approve/deny customer or HR requests.

Guiding Principles

Frame these guidelines as enabling safe innovation rather than restrictive rules. Focus on practical approaches that protect both operational data and intellectual property while maintaining usability. Remind employees which tools you provide that have appropriate protections in place (such a Team/Enterprise accounts for Claude or ChatGPT).

Example text: The intent of this policy is to enable safe and effective use of AI across [Company], not to impose unnecessary restrictions. AI is a critical tool for innovation and efficiency, and these guidelines are designed to help employees leverage AI confidently while protecting data, intellectual property, and organizational trust.

We remind employees that [Company] provides approved, enterprise-grade AI accounts (e.g., Claude Team, ChatGPT Enterprise) that have the necessary privacy, compliance, and security protections in place. Additionally, this access often provides superior functionality compared to free or personal use licenses. Use of unapproved consumer accounts is discouraged, as it may expose sensitive company or customer data.

AI should always be used to support our mission and business objectives, never in ways that create undue risk, conflict with our values, or compromise customer and regulatory obligations. This policy does not attempt to cover every possible AI use case. Instead, it provides clear standards, supported by governance processes and oversight mechanisms, to guide responsible use. For certain business functions or higher-risk applications, additional safeguards will apply, and will be further refined in the Acceptable Use Policy for each AI system. Additionally, please review the Prohibited Uses below for situations in which AI may not be used at [Company], and High-Risk Use of AI Systems below for situations in which extreme caution is required when considering using AI.

In addition, there are certain Embedded AI Tools used in existing approved [Company] software that do not require additional approval for use. For example, the use of Microsoft Word in which Microsoft

Word has embedded an AI tool to check spelling or grammar. The use of Embedded AI Tools in approved software at [Company] is permitted, provided those software tools are aligned with previous general business uses. A list of existing software tools with Embedded AI Tools that are approved at [Company] can be found here [Link].

When third-party software, services, or contractors are utilized or employed, any AI usage by software used by these parties or services must be noted and evaluated carefully. Contracted services that utilize AI technology should be considered in the same light as individual AI usage. Consult with the Legal Department about the inclusion of an AI-specific clause in any vendor or contractor agreements.

The following principles must be followed when considering using an AI system at [Company]:

- **Purpose and Scope:** AI should only be deployed for legitimate, defined business purposes that align with departmental goals. Repurposing data or outputs for unrelated objectives requires fresh approval. Any new AI tool (beyond approved embedded features in existing software) must be reviewed and approved by the AI Governance Committee. Please see the Standard Operating Procedure here [Link] for the process for getting an AI system approved. Also, see General AI Use Standards and Tool Approval sections below.

  Regulatory Consideration: *Purpose specification aligns with the EU General Data Protection Regulation (GDPR) purpose limitation principle and similar data protection laws. Any expansion of the AI's role (e.g., from workflow optimization to employee monitoring) must trigger a fresh risk assessment and approval.*

- **Subject Matter Expertise Requirement:** AI is a tool to augment expertise, not replace it. Employees must have sufficient subject matter knowledge to evaluate outputs critically. As such, individuals using an AI system must have expertise in the subject matter for which the AI is used. For example, if using AI for coding, the individual deploying the AI must have expertise in the coding language generated by the AI system.

- **Human Oversight and Accountability:** Humans remain responsible decision-makers for any AI-assisted outcome. Critical actions require explicit review and approval. The AI platform is an augmentative tool, not an autonomous decision-maker for critical actions. Human oversight is required, especially for any recommendation impacting people or significant business operations. Users must review AI recommendations and apply judgment before execution.

  Additional Consideration: *Meaningful human control aligns with the EU AI Act's anticipated requirements for human oversight in high-impact AI systems. In addition, maintaining an audit trail of AI outputs and decisions will support accountability and facilitate troubleshooting.*

- **Quality and Accuracy:** The accuracy and quality of data are paramount. Insights will only be as good as the data driving them. We will enforce strict data governance, using only relevant, up-to-date, accurate, and actively maintained datasets. Data lineage will be tracked to know how data flows into the AI and transformations applied. All AI-generated outputs (text, data, code, images, etc.) must be reviewed and refined by qualified individuals before use. Outputs should meet [Company]'s standards for accuracy, branding, and professional quality. In this sense, we view AI-generated content as a starting point, not the finished product. Furthermore, any use of an AI system must have clear objectives for the AI use as a tool and business-accepted data sets from which the AI draws. If the data sets that the AI is using are not accurate, then the information AI provides will not be accurate.

- **Bias and Fairness:** Users must evaluate AI results for bias and fairness. [Company] will perform regular testing and mitigation to ensure equitable outcomes across employees, customers, and processes. In using AI, we commit to ensuring that all AI systems treat employees, customers, and processes fairly. We will actively mitigate bias in data and algorithms to avoid discriminatory outcomes. Ethical AI requires that models do not unintentionally create adverse or unfair outcomes and avoid both overt and subtle biases.

  Regulatory Consideration: *This requirement is consistent with equal employment opportunity (EEO) laws and emerging AI fairness regulations. Bias in workforce-related tools may result in violations of anti-discrimination statutes.*

- **Data Protection:** Non-public company data must never be entered into open AI systems. When in doubt, consult the AI Governance Committee before use. Non-public [Company] information must never be put into an open AI system. If you are unsure if a system is open, please consult with the AI Governance committee and pertinent Acceptable Use Policy.

- **Documentation and Transparency:** AI usage must be documented and traceable. Tracking the use of AI is not optional and is part of your job as a Responsible AI user. This includes recording what tools are used, for what purpose, and how outputs were reviewed. Model logic, data sources, and limitations will be documented in Model Cards for transparency. Users (e.g., managers) must be informed when they are receiving an AI-generated insight and should be able to obtain an explanation in plain language. Documentation of specific AI Embedded Tools in an approved existing software tool when using that tool as intended is not required. Additionally, the use of an AI system must be documented to capture institutional knowledge. For example, if AI is used to create code and included in a larger section of code, there must be documentation as to which code section is AI-derived and who reviewed it.

  Regulatory Consideration: *Transparency is a central principle in many frameworks, including the EU AI Act's transparency obligations and various industry guidelines. This policy is designed to ensure these requirements are met by design.*

- **Contractual and Legal Compliance:** Any third-party AI use must comply with relevant contracts, licensing terms, and applicable laws. Legal review is required for inclusion of AI clauses in vendor or contractor agreements. Further, the use of an AI system must meet any terms of use or contractual limitations. Certain restrictions or terms of use may restrict [Company] use of an AI system that would otherwise be legally compliant and ethically sound

- **Standard Review Process:** Approval of an AI tool does not replace other company reviews (security, privacy, cost, HR, legal). AI systems must follow the same due diligence process as other technologies or person-centered processes. You should also ensure within your business unit that leadership is aware of the use of the AI system and has approved any use of the AI system, particularly for AI-generated content that will be relayed externally.

- **Security and Reliability:** The AI system should be robust, secure, and perform as intended. As a component of each AI system's Acceptable Use Policy, we will implement testing to ensure models meet performance targets and have fail-safes for anomalies. The AI must handle input errors or novel situations safely (e.g., if data is outside expected range, the system flags it rather than giving bad advice). Security reviews will be conducted on the AI components, just as with traditional IT systems.

Regulatory Consideration: *Ensuring the security and reliability of AI systems is a core requirement of many governance standards, such as the NIST AI Risk Management Framework, and is essential to reducing the risk of compliance breaches or harm resulting from malfunctions.*

Data and Intellectual Property

AI is only as safe as the way we use it. This section outlines how employees, contractors, and partners must handle company data, business processes, and intellectual property (IP) when interacting with AI systems. The rules vary depending on the type of AI system (open vs. closed), the industry context, and the sensitivity of the information.

Example text: When using AI tools, follow these practical guidelines to protect our information:

Safeguards for Company Data:

1. Apply the same “least privilege” principle to AI access that we use for data systems by limiting tool permissions to those who need them.
2. Always use approved enterprise AI accounts for handling sensitive or proprietary data. Personal or consumer-grade accounts must never be used with company data.
3. Treat AI outputs like a junior analyst’s work, for which you review, fact-check, and edit before sharing internally or externally. Never assume AI output is complete or correct.
4. Minimize data exposure by breaking down requests and following the minimum necessary principle. For example, instead of uploading full customer files, provide column names or synthetic examples.
5. Use generic placeholders when working with confidential matters. Keep sensitive details (names, IDs, customer records, financials) out of prompts.
6. When the task involves highly sensitive information, only use locally hosted or closed AI systems with enhanced security.
7. Report any suspected data exposure immediately. Quick escalation allows us to assess and mitigate risks effectively.

Safeguards for Intellectual Property and Proprietary Processes:

1. Never paste raw source code containing proprietary business logic into an open AI system. If code support is needed, redact or abstract sensitive logic first.
2. Use placeholder or anonymized data when describing proprietary methods or processes.
3. Avoid detailed disclosure of unique workflows in AI prompts by keeping language generic.
4. Document AI contributions to inventions or novel methods. If AI meaningfully contributes to an innovation, record it in project documentation to support future IP claims.
5. Favor secure environments for highly sensitive R&D by using only company-controlled or closed AI systems where data remains within our control.
6. Ensure all AI system terms of use and vendor contracts are reviewed for compliance with licensing, confidentiality, and IP ownership provisions (consult with legal).
7. When unsure, use your internal resources (e.g., people, process, training) to confirm whether a planned use is acceptable.

Prohibited Uses
This section should provide very clear direction on any blacklisted activities. The intention here is not to dissuade use of AI but rather ensure those operating in AI environments understand the guardrails.

Example Text: While [Company] actively encourages responsible AI adoption, there are specific uses that are strictly prohibited to protect our employees, customers, data, and reputation. These rules apply to all employees, contractors, and partners unless explicitly approved in writing by the AI Governance Committee and appropriate leadership. The following AI applications are prohibited:

- **Political or Lobbying Activity:** Using AI to draft, distribute, or otherwise support lobbying efforts directed at government officials, agencies, or entities.
- **Discrimination or Profiling:** Using AI to identify, classify, or make decisions about students, candidates, employees, contractors, or other stakeholders on the basis of protected characteristics (e.g., race, ethnicity, gender, sexual orientation, disability status, age, religion, political affiliation, etc.).
- **Exposure of Confidential or Sensitive Information:**
    - Inputting trade secrets, proprietary business logic, or confidential company data into open/public AI systems.
    - Sharing personal information (e.g., health, financial, or biometric data) into any AI tool that is not explicitly approved and secured.
    - Uploading or exposing non-public company data in prompts or training datasets.
- **Unauthorized Legal Advice:** Using AI systems to generate legal advice, draft internal or external legal documents, or provide regulatory interpretations without the involvement of the Legal Department.
- **Creation of Core Intellectual Property:** Generating inventions, algorithms, source code, or unique business methods intended for patenting, copyright, or trade secret protection without explicit legal and governance oversight.
- **Circumventing Security Controls:** Using AI to bypass or weaken company security systems, encryption, or monitoring safeguards.
- **Automated Decisions with High Human Impact:** Allowing AI systems to make final, unreviewed decisions that materially affect people (e.g., hiring, firing, compensation, medical, or disciplinary outcomes).
- **Misinformation or Manipulation:** Using AI to create, spread, or amplify misleading, false, or harmful content inside or outside the company.

These prohibitions align with obligations under data protection laws (GDPR, CCPA), anti-discrimination statutes (EEO), and emerging AI regulations (e.g., EU AI Act prohibitions on unacceptable-risk systems). Any suspected violation must be reported immediately to the AI Governance Committee and the Legal or Compliance team for investigation.

Ethical Guidelines
This section defines the principles that guide responsible AI use, emphasizing fairness, accountability, transparency, and respect for individual rights. The goal is to provide clear standards for ethical decision-making, ensuring AI is applied in line with the company's values and societal expectations.

Example text: [Company] is committed to applying AI in ways that are not only legally compliant but also ethically sound, reflecting our core values and societal expectations. There may be situations where AI use is technically permissible under law but still falls short of our ethical standards. These principles ensure that our use of AI remains transparent, fair, and aligned with the trust placed in us by employees, customers, partners, and regulators. Any use of an AI system at [Company] should conform to the following ethical guidelines:

- **Informed Consent:** Before entering personal information into any closed AI system, ensure that consent has been obtained from the individual(s) concerned. Consent must be freely given, informed, and documented where required.
- **Integrity in Use**: Users must be transparent about how AI contributed to their work. AI outputs should never be misrepresented as entirely human-authored. If AI tools are used to support tasks (e.g., drafting performance reviews, preparing reports), managers and relevant stakeholders must be made aware. Always seek approval before using AI in sensitive HR, legal, or customer-facing contexts. Additionally, you should ask permission if you desire to use an AI system tool to complete a task. For example, you should understand the Acceptable Use Policy or ask your manager and HR representative if you may use an AI system to assist in writing a performance evaluation.
- **Appropriate Content**: AI must never be used to generate material that is illegal, discriminatory, offensive, or damaging to [Company]'s brand and reputation. AI resources are company property and may not be used for harmful, harassing, or otherwise inappropriate activities.
- **Minimum Data Use**: Use only **the minimum necessary data to complete** a task. Apply data minimization, anonymization, or pseudonymization techniques wherever possible. Excessive or irrelevant data should never be included in AI training or prompts.
- **Fairness and Non-Discrimination:** AI systems must not reinforce or amplify biases. Outputs should be reviewed for fairness and equity, especially in people-related contexts (hiring, evaluation, compensation). [Company] will implement bias testing and mitigation practices consistent with EEO laws and emerging global regulations.
- **Accountability and Oversight:** Human reviewers remain accountable for all decisions informed by AI. Critical decisions may not be automated end-to-end without human oversight. Audit trails of AI inputs, outputs, and decisions must be maintained for accountability.
- **Authorized Business Use:** AI systems should not be used to perform personal tasks during work hours or on company infrastructure without explicit approval from a department leader.

High-Risk AI Systems

This section outlines how the company identifies and manages AI systems classified as high risk due to their potential impact on individuals, operations, or compliance obligations. It should clarify the additional safeguards, oversight, and approvals required before deployment. The intention is to ensure that high-risk AI is developed and used responsibly, with heightened attention to safety, accountability, and regulatory requirements.

Example text: While [Company] supports broad use of AI tools, certain applications are considered high-risk because of their potential impact on individuals' rights, company operations, or regulatory compliance. These uses require enhanced oversight, safeguards, and approvals before they can be deployed. Our approach aligns with leading global frameworks such as the EU AI Act, which defines “high-risk” systems, and the NIST AI Risk Management Framework, which recommends heightened controls.

AI systems may be classified as high-risk if they:

- Involve processing of personal or sensitive data (e.g., health, biometric, or financial data).
- Influence decisions that materially affect people’s lives or careers (e.g., hiring, promotions, compensation).
- Support safety-critical operations (e.g., in healthcare, manufacturing, or transportation).
- Impact compliance obligations, such as financial reporting, audit controls, or regulatory submissions.
- Operate in a context where bias or discrimination risk is elevated (e.g., HR, lending, housing, benefits eligibility).

Examples of High-Risk Use Cases:

Personal Data in AI Systems

- Personal data may be entered only into approved, closed AI systems with appropriate security and privacy controls.
- Use of personal data in open or consumer-grade AI systems is strictly prohibited.

Hiring and Candidate Screening

- AI may assist in candidate screening only after fairness, bias, and adverse impact testing has been conducted.
- Any AI-assisted recommendations must be reviewed by HR and hiring managers to ensure equity and compliance with EEO laws.

Personnel Decisions

- AI must not be the final authority in promotions, retention, performance evaluations, or similar HR decisions.
- Such systems must undergo bias audits, be supported by human oversight, and have clear explainability mechanisms.

Safety or Compliance-Critical Functions

- AI used in health and safety contexts (e.g., predictive maintenance for manufacturing equipment, compliance monitoring in finance) must meet robust testing and validation standards.
- A fail-safe or manual override must always be available.

Data Governance and Management

This section outlines the company's existing data governance program, highlighting the natural interplay between data governance and AI governance programs. While distinct, these programs share many similar features and may be handled in a cooperative or complimentary manner.

Example text: High-quality, well-governed data is the foundation of trustworthy AI. [Company] will apply the established Data Governance Framework (adapted for AI) to all data used in the AI platform. By enforcing these data governance practices, we aim to provide all AI services with accurate, consistent, and compliant data. This reduces the risk of errors or biases and builds confidence that any insights the AI produces are based on trustworthy data. This framework has four key components:

1. **Data Strategy and Ownership:** We maintain a clear data strategy that aligns with business objectives and the AI's purpose. Data ownership roles are defined, wherein each major dataset has an owner/steward responsible for its governance. All data initiatives for the AI should be purpose-driven and tightly aligned with the organization's vision with established data principles. Data stewards ensure day-to-day compliance with data standards, monitor data quality, and enforce policies.

2. **Data Accessibility and Lineage:** We ensure the right people have the right access to data at the right time, without compromising security. We maintain a data catalog listing all datasets feeding the AI, including metadata like source, update frequency, and any restrictions. Data lineage is tracked end-to-end: one can trace how raw data moves through transformations into the AI's inputs and outputs. This traceability aids in troubleshooting and verifying compliance (e.g., ensuring personal data isn't used outside appropriate contexts). We also encourage data democratization for transparency and discovery. Within access limits, stakeholders can inspect the data and understand its nature, fostering trust in AI outputs.

3. **Data Lifecycle Management:** Procedures will govern data through its life cycle, from ingestion and preprocessing to archival or deletion. We will integrate data management processes that, for example, handle removal of outdated or irrelevant data and prevent accumulation of stale information that could skew the AI's results. Data quality controls are applied at each stage (validation checks, outlier detection, duplication removal) to ensure the AI works with reliable information. We also implement retention schedules: data is archived or disposed of when it is no longer needed, in accordance with legal requirements. All such actions follow documented data policies and standard operating procedures.

4. **Data Privacy and Regulatory Compliance:** Given the increasing regulatory scrutiny on data, we rigorously enforce privacy and compliance measures on all AI-related data. Data is classified by sensitivity (e.g., public, internal, confidential, personal), and corresponding handling rules are applied. Sensitive personal data (PII) is subject to enhanced protections: access is limited on a need-to-know basis, data may be anonymized or masked whenever possible, and strong encryption is used both in transit and at rest. We stay updated on relevant data protection laws in all jurisdictions we operate and ensure our data practices remain compliant. For example, if the AI uses employee performance data to identify inefficiencies, we will comply with employee privacy laws.

Tool Selection and Approval

This section describes your streamlined process that enables quick adoption of newly released tools and models while maintaining oversight. Focus on enabling rather than controlling – it is always better that your people use tools within your visibility than outside of it.

Example text: At [Company], our goal is to make it easy and safe to adopt new AI tools. We recognize that employees and teams want to explore emerging models and applications quickly. By providing a streamlined approval process, we ensure adoption happens within visibility and governance rather than outside of it.

General Rule

- If an AI capability is embedded in already approved enterprise software (e.g., Microsoft 365 Copilot, Google Workspace AI, Salesforce Einstein), no additional approval is required for normal business use.
- Any other AI tool must be submitted for review and approval by the AI Governance Committee prior to use.

All new AI systems will be evaluated against the following guiding principles:

- **Lawful**: Must comply with all applicable laws, regulations, and contractual obligations. This includes licensing terms, privacy requirements (e.g., GDPR, CCPA), and sector-specific compliance rules.
- **Ethical**: Must align with [Company]'s Responsible AI principles: fairness, equity, avoidance of bias, and respect for human rights.
- **Transparent**: Clear objectives must be defined for the tool's use. Oversight and approvals should be documented and captured in the AI Registry for institutional knowledge. AI-generated content must be disclosed when legally required or when mandated by [Company] policy.
- **Necessary & Valuable**: Use of the AI system must serve a valid business purpose, such as improving efficiency, enabling innovation, or supporting strategic goals. Tools that create unnecessary costs without measurable value will not be approved.

Process for Requesting Approval

1. **Manager Pre-Approval**: The requester must first consult with their direct manager to confirm business relevance.
2. **Self-Assessment**: Ensure the tool aligns with this policy's ethical, legal, and security guidelines before submitting a request.
3. **Committee Submission**: Submit a request through the AI Approval SOP [Link] for review by the AI Governance Committee.
4. **Evaluation & Decision**: The Committee will review against the principles above, consult Legal/Compliance and Security if needed, and provide approval, rejection, or conditional approval with safeguards.
5. **Documentation**: All approved tools are logged in the AI Registry, along with permitted use cases and any restrictions.

Best Practices

- **Fast-Track Approvals**: For low-risk, low-cost tools that meet clear criteria (e.g., non-sensitive use, closed system, enterprise license available), the Committee may fast-track approval to encourage safe experimentation.
- **Pilot Programs**: New tools may be approved for limited pilots before company-wide rollout. Pilots must be documented and evaluated before expansion.
- **Vendor Assessment**: Tools from third-party vendors must undergo security and privacy reviews, including contractual checks (e.g., DPAs, terms of use) before adoption.
- **Sunset Reviews**: All approved tools will undergo periodic reviews (e.g., annually) to confirm ongoing compliance, value, and security.

AI Development and Deployment Lifecycle Overview

This section should be a complement to a fully designed standard operating procedure that specifies the design, development, and implementation (DDI) of AI systems within the company. Inclusion in this policy is meant to provide an overview of the process and not fully guide such an implementation.

Example text: All AI deployments should follow [Company] operating procedure for AI DDI. In summary, we integrate governance checkpoints and best practices at each phase of the AI solution's lifecycle, ensuring an end-to-end governed process from initial design to deployment and use. Throughout these stages, our emphasis is on integrating governance into the workflow so that it does not slow down innovation but rather accelerates safe deployment by preventing failures. By treating governance as an integral part of the AI development lifecycle (similar to quality assurance or security testing), we ensure the AI platform can scale intelligently while preserving trust and compliance.

These key stages and controls include the following:

- **Project Initiation:** Before any AI project begins, the solution owner must document the business objective and ensure it aligns with approved purposes (per the principle of purpose limitation). The AI Governance Committee must vet and approve new AI use cases via a formal intake process. A preliminary AI Risk Assessment will be conducted to identify potential ethical, legal, or operational risks of the proposed application (e.g., could the recommendations inadvertently disadvantage a group? might it rely on sensitive personal data?).

  Additional Consideration: *Tools or templates may support this approval process; for example, some organizations use an "AI intake & approval" workflow to structure proposals and review them for compliance risks.*

- **Design & Development:** During model development, teams must incorporate the core principles from the start. This includes embedding fairness (e.g., selecting training data that is representative and free of known biases or applying algorithmic techniques to correct imbalance) and privacy-by-design (minimizing use of personal data, using anonymization techniques, etc.). Developers will document design decisions and assumptions. Importantly, technical teams should work closely with compliance and domain experts at this stage. This cross-functional collaboration helps ensure ethical considerations are not overlooked. This stage should also specify the use of tools and frameworks used in bias detection, explainability, and validation tests.

  Regulatory Consideration: *In regulated domains such as finance or healthcare, AI systems must comply with sector-specific requirements (e.g., FDA guidance for medical devices, OCC guidelines for banking models). For use in the EU, the AI Act may classify certain internal systems – such as HR management or critical operations – as 'high-risk,' requiring conformity assessments, documentation, and human oversight by law.*

- **Validation & Testing:** Before an AI model or new version is rolled out, it must pass testing. This includes accuracy and performance evaluation on a test dataset, bias audits, scenario testing, and stress and security testing. A sample AI's recommendations or actions must be tested to ensure they are reasonable and do not cause inadvertent harm. Detailed records of these validation steps (test results, bias metrics, performance benchmarks) should be maintained as part of an audit trail.

Regulatory Consideration: *Maintaining audit-ready documentation and evidence of testing is increasingly expected. An ethics or compliance representative, together with the technical lead, may be required to review and formally sign off on the model's readiness.*

- **Deployment:** Deployment of the AI solution will be done in a controlled manner. Initially, the model may be launched as a pilot in a limited environment or to a subset of users to observe real-world behavior. The deployment plan should include relevant governance gates in accordance with AI core principles and data governance processes. Example questions include the following: did the model meet required accuracy? Did it undergo bias testing? Are all metadata and documentation updated? This assessment must pass and be reviewed by the AI Committee before the model is promoted to production. In accordance with standard IT policies, security reviews should be included as part of deployment. This includes activities such as scanning the AI software for vulnerabilities, ensuring API integrations have proper authentication, and that no sensitive data is exposed in logs or outputs. Finally, before full go-live, all end-users and stakeholders must receive training on the system's intended use, limitations, and their responsibilities [Link to Training]. Human oversight mechanisms must also be verified (e.g., managers know they should review suggestions, not auto-approve them). Only after satisfying all these conditions will the AI tool be considered officially deployed for broad use.

  Regulatory Consideration: *Where required – particularly in EU contexts for high-risk AI – the company should register the AI system with the relevant authorities and provide the necessary documentation (e.g., a conformity assessment or Model Card) as part of deployment.*

- **Use and Operation:** Once deployed, the AI platform's usage must adhere to established Acceptable Use Policies. Users (employees) are expected to use the AI recommendations as per training, to assist decision-making, not as absolute directives. Any usage outside the intended scope, such as using the AI to analyze data it was not designed for, is prohibited without governance approval. Where relevant, the deployment team should implement controls on inputs and outputs. All AI outputs should be logged, and significant decisions or changes influenced by AI should be traceable to either a human or AI recommendation (to maintain an audit trail of "who did what"). If at any point users feel an AI recommendation is problematic or unethical, they are empowered (and instructed) to flag it to the AI Governance Committee for review. This user feedback loop is part of operational governance to catch issues early.

Monitoring, Auditing, and Risk Mitigation

Monitoring provides the feedback loop for continuous improvement in governance processes. Avoid establishing a punitive culture where monitoring procedures are feared. Rather, these processes should be viewed as protection for AI users, the company, clients/customers, and enabling of future Responsible AI Systems.

Example Text: Governing AI is not a one-time effort. Continuous monitoring and risk management are vital once the system is in production. In summary, through vigilant monitoring and a proactive stance on risk management, we aim to catch problems early and continuously improve the AI system. This reduces the likelihood of major failures and builds confidence among stakeholders that the AI is under control. Our approach turns governance into a competitive advantage, maintaining trust while harnessing AI at scale. This policy requires proactive oversight to promptly detect issues like model performance degradation, bias, or misuse. Key measures include:

- **Performance Monitoring & Drift Detection:** AI system outputs and performance should be monitored on an ongoing basis. Monitoring should be embedded in SaaS offerings or in [Company] MLOps pipelines, using tools to log predictions and outcomes and to compare them over time. For example, if the AI suggested a process change that did not yield the expected improvement, that is feedback into the model's performance record. Drift detection techniques should be capable of flagging if the live data starts differing from training data (which could necessitate model retraining or recalibration).

- **Bias and Fairness Auditing:** Standardized and periodic fairness audits on the AI's recommendations should be included in monitoring plans. Even after careful design, biases can emerge once the system is used in practice. Built-in bias detection modules or external audit tools may be used to evaluate outcomes across different groups. The AI Governance Committee will receive reports summarizing any bias metrics or ethical concerns identified.

  Regulatory Consideration: *Regular auditing aligns with frameworks such as the NIST AI Risk Management Framework, which recommends continuous testing for harmful bias, as well as forthcoming EU requirements mandating ongoing compliance testing for high-risk AI systems.*

- **Usage Monitoring and Access Logs:** All interactions with the AI platform (inputs provided, recommendations generated, actions taken) should be logged securely. These logs serve multiple purposes: they provide an audit trail for quality investigations and help detect unintended use. We will use these to ensure the AI is being used as intended and not, for example, queried with unauthorized data. Monitoring tools can highlight anomalous usage patterns (e.g., if an employee attempts to use the AI to analyze data they normally wouldn't have access to, or if an API is being called excessively in a way that suggests a possible security breach or abuse). These records will be retained in accordance with our data retention policy (and any regulatory requirements) and reviewed periodically by compliance auditors or the AI committee.

- **Incident Response Plan:** Our goal is to remediate quickly and effectively because the only guarantee is that there will be missteps. Thus, we maintain a clear AI incident response process. If the AI system produces a faulty recommendation that leads to a significant error, or if a serious compliance issue is identified (e.g., privacy breach or discriminatory outcome), our team will react quickly. We have defined procedures similar to other IT incident responses [Link to Procedures]. Briefly, this process involves identifying and containing the issue (which may mean temporarily disabling the AI's function), analyzing root causes, remediating model deployment (e.g., roll back to a previous model version or apply a patch), and communicating to affected stakeholders. The incident response plan also includes escalation paths (when to involve legal counsel, HR, PR for communication, etc., depending on severity). We will practice drills or tabletop exercises to ensure readiness.

- **Continuous Risk Assessment:** Governance is a continuous process. We will periodically re-assess the risks of the AI platform as the business context or external environment changes.

This may be done annually or upon significant changes (like adopting a new data source, a major model upgrade, or changes in law). The AI Governance Committee will update the risk register and ensure new mitigations are implemented as needed. Part of this involves staying abreast of external developments such as new regulations, new types of AI threats, new technologies, or lessons learned from incidents at [Company] or other organizations. If regulators issue new guidance (e.g., an update to the FTC's AI guidelines or new ISO standards), we will incorporate those into our program. Our policy, therefore, is a living document.

Training and Awareness

A formal training program extends beyond governance. This section can weave the two together or focus primarily on governance controls and training. The aim is to create a culture of Responsible AI, where everyone feels responsible for upholding the policy and is encouraged to speak up if they see something questionable. An open culture is crucial, where employees feel safe to question the AI's output or the data used, without fear of retribution. This psychological safety encourages vigilance and improvement.

Example text: A critical success factor for AI governance is ensuring that all relevant personnel are educated about responsible AI use as well as maximizing the efficiencies gained by using an AI system. Technology alone cannot enforce the ethical use of AI, people must make informed decisions during development, deployment, and daily use. Investing in education and culture ensures that our workforce can effectively uphold this policy. Thus, we institute the following training and awareness measures:

- **Organization-Wide AI Governance Training:** All employees involved with the AI platform (from developers to managers and end-users) must complete mandatory training on responsible AI and this governance policy. This will be an annual training (with refreshers as needed) covering topics like understanding the AI's capabilities and limits, ethical principles (fairness, privacy, etc.), how to interpret and appropriately act on AI recommendations, and procedures for reporting issues. The training will also include basic knowledge of relevant frameworks and laws (e.g., an overview of the NIST AI Risk Management Framework, ISO/IEC 42001) and regional regulations that impact our AI use. The intention is not to make everyone a compliance expert, but to foster a baseline literacy so that teams understand key frameworks and applicable regulations that impact the business. In effect, we want a culture where employees recognize potential AI risks and their role in mitigating them.

- **Role-Specific Education:** In addition to general training, we provide tailored education for specific roles. Product and project managers will be trained on conducting AI risk assessments, documentation (e.g., creating model cards), and the process for getting approvals. Compliance and audit teams will be trained on technical aspects of AI so they can effectively audit and challenge the AI team (e.g., understanding how to interpret bias test results or drift metrics). Executives and board members will receive high-level briefings to build AI fluency, enabling them to ask the right questions and steer strategy responsibly. The company may leverage external courses or certifications for Responsible AI, as well as in-house workshops. Additionally, AI System specific training will be included to learn best practices for responsible and functional use of specific AI systems. We view training as an ongoing effort as AI technology and regulations evolve.

- **Governance in Daily Work Routines:** Beyond formal training sessions, we aim to weave AI governance awareness into daily operations. This will be achieved by providing just-in-time guidance and tools. For instance, when users interact with the AI System, there might be inline tips about interpreting results responsibly that caution that correlation does not equal causation or reminds managers to consult employees before acting on a recommendation. As a component of the ongoing Responsible AI program, regular internal communications like monthly tips or case studies will be shared via email or intranet to reinforce governance and best practices. If there are notable incidents or near-misses within the company or industry, we will anonymize and share those stories to highlight for a shared community of practice.

- **Talent and Hiring Considerations:** As we expand our AI initiatives, we will factor AI governance competencies into our hiring and talent development. As AI becomes increasingly central to operations, roles like "Responsible AI Lead" might be introduced to bring specialized expertise. For existing roles, we will incorporate accountability for AI governance to job descriptions and performance evaluations where relevant to ensure clear accountabilities and credit for contributions.

Regulatory Compliance and Legal Considerations

This section reflects the reality that, in essence, our AI governance is tightly interwoven with compliance. We treat ethical requirements and legal requirements with equal gravity as they often overlap. By being thorough in compliance, we protect our enterprise value (no costly fines or lawsuits), and by being ethical, we protect our reputation and stakeholder trust. Enterprises that rush AI deployment while sidestepping governance take on significant regulatory and reputational risk, whereas those who deploy AI responsibly will have a more sustainable advantage. This section of the policy is your commitment to be in the latter category.

Example text: This policy is designed to ensure that our AI initiatives comply with all relevant regulations, standards, and legal requirements. Because AI governance operates in a dynamic legal landscape, we commit to align AI projects with major frameworks and track evolving laws. Key regulatory and compliance considerations include:

- **Privacy Laws:** Each AI System will be evaluated for necessary compliance with data protection laws like the EU General Data Protection Regulation (GDPR) and comparable laws (CCPA/CPRA in California, LGPD in Brazil, etc.). This includes lawfully collecting and processing personal data, honoring individuals' rights, and performing Data Protection Impact Assessments (DPIAs) for AI projects that involve high-risk processing of personal data. If our AI monitors employee performance or communications, we will also adhere to applicable workplace privacy laws (and obtain employee consent or notices if required by local law). Privacy compliance is not optional. Our standard practices of data minimization and consent for data use are directly aimed at meeting these legal obligations.

- **EU AI Act and International AI Regulations:** We are monitoring the progress of the EU AI Act, which is expected to come into force soon. Should it apply to [Company] use case (the Act classifies certain AI systems as high-risk, e.g., those used in managing workers or making important decisions), we will ensure compliance with its requirements – such as conducting conformity assessments, implementing risk management, ensuring human oversight, registering the system in the EU database, etc. Even if we are not subjected to this regulation or our AI is not high-risk under the Act, we will consider its guidelines as best practices. Likewise, we will keep track of regulations in other jurisdictions: for example, any U.S. federal or state laws on AI transparency or bias (such as emerging rules on AI in hiring or credit decisions), and guidance from agencies like the FTC (which has warned against biased or deceptive AI outcomes under its consumer protection mandate). Our compliance team will update the governance committee on new rules, and this policy will be updated accordingly.

- **Sector-Specific Compliance:** If the AI platform’s recommendations touch on regulated domains, we will comply with those domain regulations. For instance, if we use AI to optimize financial processes, we will follow any relevant financial compliance (SOX controls, SEC guidance on AI use, etc.). If the AI is used in healthcare operations, we must comply with HIPAA for any patient-related data and ensure the AI’s advice does not contravene medical regulations. Our responsible AI risk assessment will explicitly check for any sector-specific compliance needs for each use case. Additionally, we align with industry standards such as SOC 2 for security and availability controls. We also look to standards like ISO/IEC 27001 (information security management) and emerging ISO standards for AI management (such as ISO/IEC 42001 on AI Management Systems once finalized) as benchmarks for our controls.

- **Standards and Frameworks:** We voluntarily adopt reputable AI governance frameworks to structure our program. A leading example is NIST’s AI Risk Management Framework, which provides a comprehensive approach to identify, mitigate, and manage AI risks. This framework will be used as a reference to ensure we cover all risk categories, such as accuracy, bias, cybersecurity, privacy, and supply chain. We also consider the OECD AI Principles as these emphasize principles like fairness, transparency, human-centered values, robustness, and accountability, all of which are reflected in our policy.

- **Record-Keeping and Reporting:** Compliance also means being ready to produce evidence of our governance. We maintain documentation of all the controls described, including data governance records, model test reports, meeting minutes from AI committee, training logs, etc. If a regulator or client requests to assess our AI’s compliance (e.g., under a contract or due diligence process), we can provide audit-ready documentation that our governance toolset helps generate. Internally, we will report major AI governance metrics to senior leadership and the board, such as number of AI systems in use, results of recent audits, and any compliance issues. This keeps governance on the radar at the highest levels of the organization.

- **Legal Agreements and Liability:** We will ensure our vendor and partner contracts account for AI-related risks. If we use third-party AI services or data, contracts will include clauses on data protection, confidentiality, and liability for issues arising from the AI. We require vendors to adhere to equivalent AI governance standards (e.g., no unlawful bias, robust security). We also clarify intellectual property and ownership of AI-created insights. From a liability perspective, we recognize that improper use of AI can lead to legal claims. By following this policy, we mitigate such risks, but we also maintain appropriate insurance and legal readiness to handle any claims. Our legal department will periodically review the AI’s functioning to ensure it is not inadvertently creating legal obligations.