MIS

POLICY | GOVERNANCE

AI Governance

Policy

SMB Compliance-Aware Template

For HIPAA, CMMC, PCI DSS, Privacy, Security, and Vendor Risk

Prepared for small to mid-sized businesses using, buying, or building AI-enabled systems.

DOCUMENT CONTROL

Organization	[Insert Company Name]
Policy Owner	[Insert Executive Sponsor / Compliance Owner]
Technical Owner	[Insert IT/Security Owner]
Effective Date	[Insert Date]
Version	1.0
Review Cadence	At least annually and after material AI, regulatory, vendor, or business changes
Applies To	Employees, contractors, vendors, service providers, systems, and business units using AI
Prepared	April 29, 2026

Use of this template

This policy is designed to be adopted by SMBs and tailored to their legal, regulatory, contractual, and customer commitments. Replace placeholders, confirm data classifications, map the policy to existing security/privacy programs, and obtain legal review before formal adoption.

MIS

TABLE OF CONTENTS

Contents

APPENDICES

1. Executive Policy Statement

[Insert Company Name] recognizes that artificial intelligence can improve productivity, security, service delivery, customer experience, analytics, and product innovation. AI also introduces material risks involving privacy, confidentiality, cybersecurity, safety, accuracy, intellectual property, bias, regulatory compliance, and customer trust.

This Policy establishes the minimum governance, security, privacy, compliance, and oversight requirements for the evaluation, approval, use, development, procurement, and monitoring of AI systems. The Company will use AI only in a manner that is lawful, ethical, secure, transparent where appropriate, accountable, and consistent with customer contracts and regulatory obligations.

AI systems may not be used to bypass existing security policies, data handling rules, change management, vendor due diligence, incident response, legal review, or compliance requirements. Where this Policy conflicts with a stricter legal, regulatory, contractual, or customer requirement, the stricter requirement controls.

2. Purpose and Objectives

The purpose of this Policy is to provide a practical governance framework for AI adoption by small to mid-sized businesses, including organizations subject to regulated-data obligations such as HIPAA, CMMC, PCI DSS, GLBA, state privacy laws, GDPR, and customer-specific security requirements.

- Create a repeatable process for AI use-case intake, risk assessment, approval, and monitoring.
- Prevent unauthorized disclosure of regulated, confidential, proprietary, or customer data through AI tools.
- Ensure AI vendors and embedded AI capabilities are reviewed before production use.
- Establish human oversight expectations for AI-generated outputs and automated actions.
- Align AI governance with existing cybersecurity, privacy, compliance, vendor risk, incident response, and change management programs.
- Maintain auditable evidence of AI decisions, approvals, exceptions, incidents, and periodic reviews.

3. Scope

This Policy applies to all AI systems used, purchased, configured, integrated, or developed by or on behalf of the Company, including internally used tools, customer-facing tools, AI-enabled SaaS applications, copilots, chatbots, virtual agents, machine learning models, generative AI, predictive

analytics, automated decision systems, AI-assisted cybersecurity tools, and AI features embedded in existing products.

- All employees, contractors, temporary workers, consultants, interns, and service providers using Company systems or Company data.
- All Company data, customer data, employee data, regulated data, source code, system configurations, business records, and intellectual property.
- All environments, including production, development, testing, sandbox, cloud, endpoint, mobile, and third-party platforms.
- AI tools used with Company accounts, personal accounts for Company work, browser plugins, APIs, automation tools, scripts, extensions, and integrations.

Out of scope

Personal use of AI that does not involve Company resources, Company data, customer data, regulated data, Company intellectual property, or Company business activities. Personal use remains subject to law and Company conduct policies when it affects the Company.

4. Definitions

Term	Definition
AI System	Software, model, service, workflow, embedded feature, or automated process that uses artificial intelligence, machine learning, generative AI, natural language processing, computer vision, autonomous agents, or predictive analytics to produce outputs, recommendations, decisions, actions, or content.
Generative AI	AI that creates or transforms text, code, images, audio, video, designs, analyses, summaries, or other content based on prompts, context, training data, or retrieval sources.
AI Use Case	A business process or activity in which an AI system is used to assist, augment, automate, recommend, decide, generate, classify, summarize, detect, or act.
AI Owner	The business owner accountable for the AI use case, its intended purpose, approval, performance, monitoring, and compliance.
Regulated Data	Information subject to legal, regulatory, contractual, or industry-specific controls, including ePHI, CUI, FCI, cardholder data, sensitive authentication data, PII, NPI,

Term	Definition
	biometric data, employee data, export-controlled data, and customer confidential data.
ePHI	Electronic protected health information subject to HIPAA Security Rule obligations when created, received, maintained, or transmitted by covered entities or business associates.
CUI / FCI	Controlled Unclassified Information and Federal Contract Information requiring protection under federal contract obligations, including CMMC and related NIST security requirements where applicable.
Cardholder Data	Payment card data, including PAN and other PCI DSS-scoped data. Sensitive authentication data includes full track data, CAV2/CVC2/CVV2/CID, PINs, and PIN blocks.
High-Impact Decision	An AI-supported decision or recommendation that may materially affect employment, credit, housing, healthcare, insurance, education, legal rights, safety, financial eligibility, access to services, or contractual obligations.
Human Oversight	Meaningful review by an accountable person with enough authority, context, and time to challenge, verify, approve, reject, or override AI output before it is relied upon.

5. Governance Principles

Principle	Policy Intent
Accountability	Every AI use case must have a named business owner, technical owner, and approval path proportional to risk.
Data minimization	Only the minimum necessary data may be used. Regulated or confidential data must not be entered into unapproved AI systems.
Security by design	AI systems must follow access control, encryption, logging, vulnerability management, secure development, and incident response requirements.
Human oversight	AI outputs are advisory unless explicitly approved for automation. Humans remain accountable for decisions and customer-facing actions.
Transparency and notice	Users, customers, and affected individuals should receive notice of AI use when required by law, contract, risk, or customer expectations.
Fairness and non-discrimination	AI must not be used in a way that unlawfully discriminates, creates unfair outcomes, or relies on unsupported bias claims.

Principle	Policy Intent
Accuracy and reliability	AI outputs must be validated before use, especially in regulated, customer-facing, security, financial, healthcare, or contractual contexts.
Vendor accountability	AI vendors must be assessed for security, privacy, compliance, contract terms, data use, retention, subprocessors, and incident commitments.
Auditability	The Company must retain sufficient evidence to demonstrate approvals, risk reviews, control operation, incidents, training, and periodic reassessment.

6. Roles and Responsibilities

For SMBs, the AI Governance Committee may be a small cross-functional group, or the responsibilities may be assigned to an existing security, compliance, or risk committee. At minimum, there must be a named executive sponsor, security owner, and business owner for each approved AI use case.

Role	Responsibilities
Executive Sponsor	Approves policy adoption, risk appetite, high-risk AI use cases, funding, and unresolved exceptions.
AI Governance Committee	Reviews moderate and high-risk use cases, approves standards, tracks AI inventory, reviews incidents, and confirms regulatory alignment.
AI Owner / Business Owner	Defines business purpose, data used, expected outputs, success criteria, user group, human oversight, and ongoing monitoring.
IT / Security	Performs security reviews, access control, logging, network/API security, vulnerability management, data loss prevention, and incident support.
Privacy / Compliance	Evaluates regulated data, privacy notices, consent, DPAs, BAAs, PCI/CMMC/HIPAA implications, and evidence requirements.
Legal / Contract Owner	Reviews customer contracts, vendor terms, IP ownership, indemnities, confidentiality, data processing, and regulatory risk.
HR / People Operations	Approves workforce-impacting AI uses, employee monitoring concerns, hiring or performance-related tools, and employee training.
System Administrator	Configures approved AI tools, SSO/MFA, role-based access, retention settings, opt-out from model training, audit logs, and deprovisioning.

Role	Responsibilities
Users	Use only approved tools, follow data rules, verify outputs, label AI involvement when required, and report incidents or suspected misuse.

6.1 RACI Guide

R = Responsible, A = Accountable, C = Consulted, I = Informed.

Activity	Executive	AI Owner	IT/Sec	Priv/Comp	Legal	Users
Approve AI policy	A	C	R	C	C	I
Approve low-risk use case	I	A/R	C	C	I	I
Approve moderate-risk use case	C	A/R	R	R	C	I
Approve high-risk or regulated-data use case	A	R	R	R	R	I
Vendor due diligence	I	R	A/R	R	C	I
AI incident response	A/I	R	R	R	R	R

7. AI System Inventory and Classification

The Company must maintain an AI System Inventory covering approved, pending, retired, and denied AI use cases. Unlisted AI systems are not approved for Company use unless explicitly categorized as pre-approved low-risk tools by IT/Security and Compliance.

- AI system name, vendor, version, model, embedded feature, or integration.
- Business owner, technical owner, administrator, user population, and business process.
- Data classifications used, including whether ePHI, CUI, FCI, cardholder data, PII, confidential client data, employee data, source code, or secrets are involved.
- Deployment model: public SaaS, enterprise SaaS, private cloud, on-premises, API, embedded feature, browser extension, agent, or internally developed model.

- Vendor contractual safeguards, security attestations, BAAs, DPAs, PCI AOCs, CMMC/FedRAMP relevance, subprocessors, data residency, retention, and deletion capabilities.
- Approved use, prohibited use, risk tier, approval date, review date, residual risk, open issues, monitoring requirements, and retirement criteria.

Data Classification and AI Handling

Classification	Examples	AI Handling Rule
Public	Approved public marketing content, public documentation, public websites, public regulations	Permitted in approved AI tools. Verify outputs before publishing.
Internal	Internal procedures, non-sensitive project notes, generic operational information	Permitted in approved enterprise AI tools when no confidential or regulated data is included.
Confidential	Customer names, contracts, pricing, internal financials, non-public strategy, proprietary source code	Requires approved tool, access control, vendor review, and data owner approval.
Regulated Restricted	ePHI, CUI, FCI, cardholder data, SAD, PII, NPI, biometric data, HR records, credentials, secrets, incident details	May be used only in specifically approved systems with documented legal, compliance, security, and contractual controls.

8. AI Risk Tiering and Approval Requirements

Each AI use case must be assigned a risk tier before production use. Risk tiering must consider data sensitivity, user population, customer impact, autonomy, legal/regulatory exposure, security impact, vendor maturity, output reliance, and potential harm if the AI is wrong, biased, unavailable, manipulated, or compromised.

Risk Tier	Examples	Minimum Approval
Tier 0 — Prohibited	Unapproved AI tools using regulated data; AI used for impersonation, deception, malware, credential theft, unauthorized surveillance, unlawful discrimination, or unsupported	Not allowed unless reclassified by Executive Sponsor, Legal, Security, and Compliance after documented change in facts.

Risk Tier	Examples	Minimum Approval
	autonomous decisions affecting rights/safety.	
Tier 1 — Low	Drafting public content, summarizing public documents, brainstorming, formatting, grammar, generic code examples with no secrets or proprietary data.	Approved tool list; user training; manager oversight as needed.
Tier 2 — Moderate	Internal productivity using confidential data, internal analytics, ticket summaries, customer-support drafts, sales enablement, internal knowledge search.	AI Owner approval; IT/Security review; data owner approval; inventory entry; annual review.
Tier 3 — High	Customer-facing AI; AI using regulated or customer confidential data; cybersecurity automation; RAG over sensitive repositories; HR, finance, legal, or healthcare support; external communications at scale.	AI Governance Committee approval; privacy/compliance review; vendor review; testing; human oversight; monitoring; quarterly review.
Tier 4 — Critical / Regulated Enterprise	AI processing ePHI, CUI, FCI, cardholder data, or other regulated data; AI that triggers actions in production systems; high-impact decisions; autonomous agents with material authority.	Executive approval; legal/compliance signoff; documented risk assessment; contractual safeguards; change management; pre-production testing; logging; incident playbook; review at least quarterly.

Scenario-Based Tier Defaults

Scenario	Default Tier	Minimum Approval	Minimum Review
Public data only	Tier 1	Approved AI tool; trained user	Annual
Internal data	Tier 1 or 2	AI Owner; IT/Security if tool not pre-approved	Annual
Confidential company or customer data	Tier 2 or 3	AI Owner; Data Owner; IT/Security; Compliance if customer commitments apply	Annual or quarterly if customer-facing
ePHI / HIPAA	Tier 4	Executive; Privacy/Compliance; IT/Security; Legal; BAA review	Quarterly

Scenario	Default Tier	Minimum Approval	Minimum Review
CUI / FCI / CMMC	Tier 4	Executive; CMMC Program Owner; IT/Security; Legal; supplier flowdown review	Quarterly or per contract
Cardholder data / PCI	Tier 4	Executive; PCI Owner; IT/Security; Legal; QSA/assessor input where applicable	Quarterly
High-impact decisions	Tier 3 or 4	Executive; Legal; HR/Compliance; documented human oversight	Quarterly
Autonomous agents with system actions	Tier 3 or 4	IT/Security; Change Advisory; AI Governance Committee; Executive for material actions	Quarterly

9. Acceptable and Prohibited Uses

9.1 Acceptable Uses

- Using approved AI tools to draft, summarize, reformat, translate, or improve non-sensitive content, subject to human review.
- Generating internal training materials, knowledge-base drafts, scripts, code examples, or workflow suggestions that do not contain confidential, regulated, or customer data unless separately approved.
- Assisting with security analysis, log triage, documentation, query generation, or reporting when data handling rules and access controls are followed.
- Using AI to support customer service, sales, finance, operations, or delivery processes when the use case is inventoried, approved, monitored, and subject to appropriate human oversight.
- Using approved enterprise AI tools configured to prevent vendor training on Company data, limit retention, enforce SSO/MFA, log use, and support deletion/export where feasible.

9.2 Prohibited Uses

- Entering, uploading, pasting, or exposing regulated data into any public, personal, trial, freemium, consumer, browser-extension, or unapproved AI tool.
- Entering credentials, API keys, private keys, passwords, tokens, secrets, seed phrases, sensitive incident details, unreleased vulnerabilities, or security architecture into AI systems not approved for such data.

- Using AI to make final decisions involving employment, compensation, credit, healthcare, insurance, housing, legal rights, safety, or access to services without documented human oversight and legal/compliance approval.
- Generating or sending customer, legal, medical, financial, regulatory, or contractual advice without review by an authorized professional or business owner.
- Using AI to impersonate individuals, create deceptive content, create fake reviews, generate unauthorized deepfakes, misrepresent Company identity, or make unsupported claims about AI capabilities.
- Using AI to develop malware, phishing, social engineering, credential theft, evasion, vulnerability exploitation against unauthorized systems, or other malicious activity.
- Using AI-generated content, code, or designs without reviewing for intellectual property, license, security, privacy, accuracy, and confidentiality concerns.
- Connecting AI agents to production systems, payment systems, administrative consoles, ticketing systems, email sending, customer data stores, or code repositories without change management and security approval.

10. Data Protection and Privacy Requirements

AI use must comply with the Company data classification standard, acceptable use policy, privacy policy, information security policy, customer agreements, and regulatory obligations. AI systems must be designed and configured to minimize data exposure and prevent unauthorized disclosure.

- Use approved AI tools only. Public or personal AI accounts are not approved for Company confidential or regulated data.
- Apply minimum necessary principles. Redact, anonymize, aggregate, or tokenize data before AI use whenever feasible.
- Do not allow vendors to train public or shared models on Company, customer, employee, or regulated data unless Legal, Privacy, Compliance, and Executive Sponsor explicitly approve the arrangement in writing.
- Require SSO, MFA, role-based access control, least privilege, access reviews, and timely deprovisioning for enterprise AI systems.
- Require encryption in transit and at rest for AI systems that store, process, transmit, or retrieve confidential or regulated data.
- Ensure audit logs are available for access, prompt activity, file uploads, data retrieval, administrative changes, integrations, and material automated actions, when supported by the system.

- Configure retention periods to the shortest operationally necessary timeframe and align deletion with legal, regulatory, contractual, and evidence requirements.
- Perform privacy review when AI processes PII, employee data, customer data, biometric data, location data, health data, financial data, or data from minors.
- Review cross-border data transfer, data residency, subprocessors, and cloud hosting locations for AI systems that process regulated or customer data.
- Do not use AI to evade consent, notice, purpose limitation, contractual restrictions, or privacy commitments made to customers, employees, or consumers.

Data Type Handling Rules

Data Type	AI Handling Rule
Credentials / secrets	Never enter into AI prompts or upload to AI tools. Rotate immediately if exposed.
Source code	Allowed only in approved tools. Remove secrets, proprietary client code, unreleased vulnerabilities, and customer identifiers unless approved.
Customer contracts / pricing	Approved enterprise tools only; data owner approval required; no external disclosure.
ePHI	Approved HIPAA-scoped systems only; BAA required for vendor processing; minimum necessary; audit logging and access control required.
CUI / FCI	CMMC-scoped environment only; AI vendor and integrations must be assessed against applicable contractual requirements.
Cardholder data	Do not enter into AI prompts unless the AI system is explicitly included in PCI scope and approved. Sensitive authentication data must never be retained after authorization.
Employee / HR data	HR and Legal review required for AI uses involving hiring, monitoring, performance, compensation, or termination.

11. Compliance-Specific Requirements

This section applies when the Company is subject to the listed framework by law, contract, customer requirement, assessment scope, or business commitment. It is designed as a control overlay; it does not replace the underlying compliance program.

Framework / Domain	Minimum AI Governance Requirements
HIPAA / ePHI	AI tools that create, receive, maintain, or transmit ePHI must be approved for HIPAA use, covered by an appropriate BAA when a vendor acts as a business associate, configured with access controls, audit controls, encryption where appropriate, minimum necessary use, and incident/breach escalation. AI-generated patient or clinical content must receive qualified human review before use.
CMMC / CUI / FCI	AI systems processing, storing, or transmitting FCI or CUI must be within the CMMC-scoped system boundary or otherwise authorized through a documented compliance architecture. Supplier flowdowns, current CMMC status where required, SSP/POA&M evidence, access controls, logging, encryption, and NIST SP 800-171-aligned safeguards must be addressed before use.
PCI DSS / Cardholder Data	AI systems must not receive cardholder data unless they are explicitly approved and included in the PCI scope. PAN must be masked or tokenized where possible. Sensitive authentication data must not be stored after authorization. AI vendors touching cardholder data must be reviewed as service providers and provide applicable PCI evidence.
GLBA / Financial Customer Information	AI use involving nonpublic personal information must follow safeguards, vendor oversight, least privilege, retention, and consumer privacy commitments.
GDPR / EU AI Act / International Privacy	AI use involving EU/UK personal data must address lawful basis, transparency, purpose limitation, data minimization, data subject rights, cross-border transfer, automated decision-making, and AI-specific obligations where applicable.
State Privacy Laws	AI use involving personal data must support consumer rights, disclosures, sensitive data restrictions, opt-outs, data processing agreements, and profiling/automated-decision requirements where applicable.
SOC 2 / ISO 27001 / ISO 42001	AI governance evidence should align with existing control systems: risk assessment, access control, change management, vendor risk, incident response, logging, policy training, and management review.
FTC / Advertising / Consumer Protection	Company claims about AI capabilities, accuracy, safety, bias, compliance, or business outcomes must be truthful, substantiated, and not deceptive or unfair. Marketing must not overstate AI functionality or imply human review where none exists.
Employment / HR Compliance	AI tools used in hiring, screening, monitoring, performance management, compensation, scheduling, or termination must receive HR, Legal, and Compliance approval and be reviewed for fairness, transparency, accessibility, and applicable notice requirements.

12. Vendor, SaaS, and Third-Party AI Requirements

AI vendors and AI-enabled features must be reviewed before use when they process Company, customer, confidential, employee, or regulated data; integrate with Company systems; produce customer-facing outputs; make recommendations affecting people; or perform automated actions.

- Vendor terms must not permit use of Company, customer, employee, or regulated data for vendor model training unless explicitly approved.
- Contracts must address confidentiality, data ownership, permitted processing, retention, deletion, export, subprocessors, data location, breach notification, audit rights, compliance obligations, service levels, indemnity, and termination assistance.
- For ePHI, confirm BAA requirements before processing PHI. For cardholder data, confirm PCI service-provider responsibilities and evidence. For CUI/FCI, confirm contract flowdowns and CMMC/NIST responsibilities.
- Evaluate vendor security documentation such as SOC 2 Type II, ISO 27001, pen test summary, vulnerability management, secure SDLC, encryption, SSO/MFA, RBAC, logging, availability, and incident response.
- Assess AI-specific controls, including model training data use, retention of prompts and outputs, isolation between customers, RAG access controls, prompt injection protections, evaluation methods, monitoring, explainability where feasible, and change notification for model updates.
- Do not enable embedded AI features by default for regulated or confidential data stores until the feature is reviewed and approved.

Required Vendor Review Areas

Review Area	Required Questions
Data use	What data is processed? Is data used for training? Can training be disabled? Are prompts/outputs retained?
Security	SSO/MFA, RBAC, encryption, logging, isolation, vulnerability management, security certifications, pen test evidence.
Compliance	BAA, DPA, PCI AOC, SOC 2, ISO 27001, FedRAMP, CMMC relevance, regulatory obligations, customer flowdowns.
Privacy	Data residency, subprocessors, international transfers, retention/deletion, data subject rights, employee/consumer notice.

Review Area	Required Questions
AI risk	Accuracy testing, bias testing, hallucination controls, content filtering, prompt injection defense, model update notification.
Operational	SLA, support, incident notification, audit support, export, backup, portability, termination, business continuity.

13. AI Development, Integration, and Deployment Controls

Internally developed AI systems, custom AI integrations, RAG solutions, agents, model fine-tuning, AI-enabled automations, and API-based AI workflows must follow secure development, change management, privacy, and compliance requirements before production use.

- Document the intended purpose, users, data sources, model/provider, system architecture, integrations, access controls, outputs, limitations, and failure modes.
- Perform threat modeling for data leakage, prompt injection, indirect prompt injection, jailbreaks, insecure plugins, excessive agency, model inversion, unauthorized retrieval, insecure output handling, and supply-chain risk.
- Use least privilege for AI service accounts, API keys, retrieval connectors, automation accounts, and agents. Secrets must be stored in approved secret management systems.
- Apply change management for production deployment, model changes, prompt/template changes, retrieval-source changes, agent tool changes, and material vendor or model updates.
- For RAG systems, enforce source-level access control. Users must not retrieve data through AI that they could not access directly through approved systems.
- For AI agents, require defined action boundaries, transaction limits, approval gates for material actions, sandbox testing, emergency stop/disable capability, monitoring, and rollback procedures.
- Evaluate outputs using representative test cases before launch and after material changes. High-risk systems must include documented acceptance criteria and known limitations.
- Maintain a model card, system card, or use-case record for high-risk AI systems covering purpose, data, limitations, testing, monitoring, owners, and residual risks.
- Do not deploy AI that can send external communications, modify production systems, approve transactions, change access rights, or execute code without documented approvals and controls.

AI Agent Control Requirements

Agent Control	Requirement
Action authority	Define exactly what the agent can do, where, and under whose authority.
Least privilege	Give the agent only required permissions; no standing admin rights unless separately approved.
Human approval	Require approval for payments, customer notices, account changes, production changes, deletions, or legal/compliance actions.
Kill switch	Provide a documented method to immediately disable agent access and integrations.
Logging	Log prompts, retrieved context, decisions, tool calls, user approvals, and actions where technically feasible.
Testing	Test in a sandbox with adversarial prompts, malformed inputs, and failure conditions before production.

14. AI Output Review, Human Oversight, and Quality Assurance

AI outputs are not authoritative unless specifically validated. Users remain responsible for confirming accuracy, completeness, appropriateness, legality, confidentiality, and security before relying on AI output.

- Verify factual claims, citations, calculations, legal/regulatory statements, medical/clinical statements, financial recommendations, security guidance, and customer-impacting outputs using authoritative sources.
- Review AI-generated code for security vulnerabilities, licensing concerns, correctness, secrets exposure, and compliance with secure coding standards.
- Customer-facing AI content must be reviewed and approved according to the risk tier and applicable communications process.
- High-impact decisions require meaningful human review and documented rationale. A human reviewer must have authority to override the AI recommendation.
- AI-generated legal, healthcare, tax, financial, compliance, or cybersecurity advice must not be sent externally or used operationally without qualified review.
- AI-generated content should be labeled or disclosed when required by law, contract, customer expectations, or Company communications standards.

- Outputs that appear biased, discriminatory, unsafe, hallucinated, malicious, privacy-invasive, or inconsistent with Company policy must be reported to the AI Owner or IT/Security.

15. Security Monitoring, Logging, and Audit Evidence

The Company must maintain monitoring and evidence appropriate to each AI use case's risk tier. Evidence should be sufficient to demonstrate governance, security, privacy, vendor oversight, and compliance activities during audits, assessments, customer due diligence, and incident investigations.

- Maintain AI inventory entries, risk assessments, approval records, vendor reviews, contracts, data processing terms, BAAs/DPAs/AOCs where applicable, and change records.
- Log user access, administrative changes, file uploads, API integrations, material prompts/outputs, retrieval events, and automated actions where supported and legally permissible.
- Conduct access reviews at least annually for low/moderate-risk systems and quarterly for high-risk or regulated-data systems.
- Review vendor security/compliance evidence at least annually and upon material change, incident, contract renewal, or expansion of data scope.
- Monitor for data leakage, unusual prompt activity, excessive data uploads, unapproved AI tools, shadow AI, anomalous API use, and unauthorized integrations.
- High-risk systems must have documented performance, accuracy, safety, and control monitoring. Results must be reviewed by the AI Owner and Governance Committee.

Suggested AI Governance Metrics

Metric	Example Measurement
Inventory completeness	% of known AI systems recorded and risk-tiered
Approved-tool adoption	% of AI usage occurring in approved enterprise tools
Training completion	% of workforce completing AI acceptable use training
Vendor review coverage	% of AI vendors with current security/privacy review
Incident rate	AI-related incidents, near misses, data exposure events, and policy violations

Metric	Example Measurement
High-risk review status	% of high-risk systems with current quarterly review and monitoring evidence
Output quality	Error trends, hallucination reports, rejected outputs, customer complaints, and corrective actions

16. Incident Response and Regulatory Escalation

AI-related incidents must be reported and handled through the Company incident response process. AI incidents may trigger contractual notice, privacy breach analysis, HIPAA breach assessment, PCI incident procedures, CMMC/DoD reporting obligations, employment review, customer notice, or legal privilege considerations.

- Regulated, customer, employee, confidential, or proprietary data entered into or exposed through an unapproved AI tool.
- AI vendor breach, prompt/output exposure, cross-tenant data leakage, unauthorized model training, or unexpected retention of Company data.
- AI output causes or could cause customer harm, financial loss, legal/regulatory exposure, discrimination, unsafe action, or production outage.
- AI agent performs an unauthorized action, sends unauthorized communications, changes access rights, executes code, or modifies production systems unexpectedly.
- Prompt injection, indirect prompt injection, malicious retrieval, data exfiltration, jailbreak, or manipulation of AI output is detected.
- AI-generated code or automation creates a security vulnerability, licensing issue, data exposure, or operational disruption.

16.1 AI Incident Response Steps

- Report immediately to IT/Security, the AI Owner, and the incident response lead using the approved incident channel.
- Preserve relevant evidence, including user, prompt, output, file upload, retrieval, API, tool-call, administrative, and vendor logs where available.
- Contain the incident by disabling access, revoking tokens, rotating exposed secrets, pausing integrations, blocking tools, or taking the AI system offline.
- Assess data types involved, affected individuals/customers, systems impacted, contractual commitments, and regulatory reporting obligations.

- Engage Legal, Privacy/Compliance, HR, customer account leadership, cyber insurance, outside counsel, or forensics as required by severity.
- Remediate root cause, update controls, notify affected parties where required, document lessons learned, and determine whether the use case remains approved.

17. Training, Awareness, and Workforce Rules

All workforce members must receive AI acceptable use training before using Company-approved AI tools and at least annually thereafter. Additional training is required for AI owners, developers, administrators, customer-facing users, and users handling regulated data.

- Approved vs. prohibited AI tools and use cases.
- Data classification rules and examples of regulated data that must not be entered into unapproved AI systems.
- How to verify AI outputs and avoid hallucination, bias, unsafe advice, and overreliance.
- Prompting rules, confidentiality expectations, customer contract obligations, and intellectual property considerations.
- Security risks, including prompt injection, data leakage, malicious code, secrets exposure, phishing, impersonation, and unsafe agents.
- Incident reporting requirements and examples of reportable AI events.

18. Records, Retention, Exceptions, and Enforcement

18.1 Records and Retention

- Retain AI inventory records, risk assessments, approvals, vendor reviews, contracts, BAAs/DPAs/AOCs, security evidence, testing records, incident records, training records, and exceptions according to the Company retention schedule.
- Prompt and output logs must be retained only as necessary for security, compliance, quality, and operational purposes and must be protected according to data classification.
- Records involving regulated data must be retained, protected, and disposed of according to applicable law, contract, and evidence requirements.

18.2 Exceptions

- Exceptions must be documented, risk-assessed, approved by the appropriate authority, time-bound, and reviewed at expiration.

- Exceptions involving regulated data, high-impact decisions, production access, customer commitments, or legal obligations require Executive Sponsor, IT/Security, Compliance, and Legal approval.
- Emergency exceptions must be documented retroactively within five business days and reviewed for corrective action.

18.3 Enforcement

Violations of this Policy may result in removal of AI access, disciplinary action, contract remedies, customer notification, regulator notification, legal action, or termination of employment or vendor relationship, depending on severity and applicable law or contract.

19. Review and Continuous Improvement

This Policy must be reviewed at least annually and after material changes to AI technology, business operations, regulatory obligations, customer commitments, cyber threats, vendor practices, or incidents. The AI Governance Committee or assigned policy owner must track improvement actions to closure.

- Quarterly review of high-risk AI systems and regulated-data use cases.
- Annual review of approved AI tools, vendors, access, training, incident trends, and exceptions.
- Review after major vendor model changes, new embedded AI features, new regulatory requirements, new customer contracts, mergers/acquisitions, or significant incidents.
- Update the AI inventory and approval decisions when business purpose, data type, user group, integration, vendor, model, geography, output use, or automation level changes.

APPENDICES

Reference Materials & Forms

Appendix A — AI Use Case Intake Form

Complete this intake before using a new AI tool, enabling an AI feature, integrating an AI API, deploying an AI workflow, or expanding an existing AI use case to new data or users.

Field	Required Information
Requestor	Name, department, manager, email
AI Owner	Business owner accountable for use case
Technical Owner	System admin, developer, or IT owner
AI system / vendor	Product name, model, version, hosting, URL, contract owner
Business purpose	Problem solved, expected benefits, users, outputs, decisions/actions supported
Data used	Data classification, source systems, regulated data, customer data, employee data, code, files
Output use	Internal draft, customer-facing content, recommendation, decision support, automated action
Autonomy	Human-in-the-loop, human-on-the-loop, fully automated, agent with tool access
Integrations	Systems connected, API permissions, identity/access model, file stores, email, ticketing
Risk tier proposed	Tier 1, 2, 3, or 4 with rationale
Compliance triggers	HIPAA, CMMC, PCI, GLBA, GDPR, state privacy, employment, customer contracts
Vendor status	New, existing, embedded feature, approved, pending review, denied
Testing plan	Accuracy, security, privacy, prompt injection, bias/fairness, failover, rollback
Monitoring plan	Metrics, logs, access reviews, output review, incident triggers, review cadence
Approvals	AI Owner, IT/Security, Compliance, Legal, Executive, date

Appendix B — AI Risk Assessment Checklist

Data and privacy

- ☐ Does the use case process regulated data, customer confidential data, employee data, biometric data, location data, or data from minors?
- ☐ Can the data be redacted, minimized, aggregated, de-identified, or tokenized before AI processing?
- ☐ Are notices, consents, privacy policies, DPAs, BAAs, or cross-border transfer controls required?
- ☐ Does the vendor use prompts, outputs, files, embeddings, or retrieved data to train models?

Security

- ☐ Does the AI system support SSO, MFA, RBAC, logging, encryption, retention control, deletion, and access review?
- ☐ Could prompt injection, malicious files, insecure plugins, or retrieval abuse expose data or cause unauthorized actions?
- ☐ Are secrets, credentials, API keys, source code, vulnerabilities, or architecture details involved?
- ☐ Are AI agents or automations limited, monitored, and reversible?

Compliance and legal

- ☐ Does the use case involve HIPAA, CMMC, PCI DSS, GLBA, GDPR, state privacy, employment law, consumer protection, or customer contract commitments?
- ☐ Are AI claims, marketing claims, accuracy claims, or compliance claims substantiated?
- ☐ Does the use case create intellectual property, copyright, licensing, open-source, confidentiality, or privilege concerns?
- ☐ Would a customer, regulator, auditor, assessor, or court expect documented human oversight?

Operational and quality

- ☐ What is the consequence if the AI output is wrong, biased, unavailable, manipulated, or incomplete?
- ☐ How will outputs be verified before use?
- ☐ Are model limitations documented and communicated to users?
- ☐ Are post-deployment metrics, issue reporting, rollback, and periodic review defined?

Appendix C — Approved AI Tool Register

Tool name	Vendor	Approved data	Approved users	Risk tier	Owner	Review due
[Example] Enterprise AI Assistant	[Vendor]	Public/Internal only	All trained employees	Tier 1	IT	[Date]
[Example] Support Ticket Summarizer	[Vendor]	Customer confidential, no regulated data	Support team	Tier 2	Support Director	[Date]
[Example] HIPAA AI Documentation Tool	[Vendor]	ePHI approved with BAA	Clinical ops	Tier 4	Compliance Officer	[Date]

Appendix D — Vendor Due Diligence Questionnaire

- ☐ Describe all data processed, stored, transmitted, embedded, indexed, logged, retained, or used for training.
- ☐ Confirm whether customer data, prompts, outputs, files, embeddings, or retrieval content are used to train, tune, evaluate, or improve models.
- ☐ Provide data retention, deletion, export, legal hold, backup, and termination procedures.
- ☐ List subprocessors, hosting regions, data residency options, and cross-border transfer mechanisms.
- ☐ Provide SOC 2 Type II, ISO 27001, PCI AOC, FedRAMP, CMMC relevance, BAA, DPA, penetration test summary, or other evidence as applicable.
- ☐ Describe SSO, MFA, RBAC, SCIM/deprovisioning, audit logging, encryption, customer-managed keys, and administrative controls.
- ☐ Describe AI-specific safeguards for prompt injection, data leakage, hallucination, harmful output, bias, unsafe instructions, and model changes.
- ☐ Provide incident notification commitments, support SLAs, vulnerability disclosure process, and breach cooperation terms.
- ☐ Explain how customer data is isolated from other customers and how retrieval permissions are enforced.
- ☐ Identify whether the tool includes autonomous agents, plugins, connectors, external tool calls, or access to email, file systems, code repositories, ticketing, payment, CRM, or production systems.

Appendix E — Workforce AI Quick Rules

Distribute these rules to every workforce member as part of AI acceptable use training and onboarding.

Rule Type	Rule
Do	Use approved Company AI tools only for Company work.
Do	Assume AI outputs can be wrong. Verify before relying on them.
Do	Redact customer, employee, regulated, confidential, and security-sensitive details unless the use case is approved for that data.
Do	Report accidental disclosure, suspicious outputs, unapproved AI use, or AI incidents immediately.
Do not	Enter PHI/ePHI, CUI, FCI, cardholder data, employee records, customer confidential data, credentials, secrets, or source code into unapproved AI tools.
Do not	Use personal AI accounts, browser extensions, or freemium AI tools for Company data.
Do not	Send AI-generated legal, medical, financial, compliance, security, or customer-impacting content without qualified review.
Do not	Let AI agents act in production systems, send external messages, modify records, approve transactions, or change access without approval.

Appendix F — Compliance Control Mapping

Control Area	Policy Evidence	Framework Alignment
AI inventory	Vendor/product/use-case register; owner; risk tier; data scope; review date	NIST AI RMF, ISO 42001, NIST CSF, SOC 2, ISO 27001, CMMC
Data classification	Rules for Public, Internal, Confidential, Regulated Restricted data	HIPAA, CMMC, PCI DSS, GLBA, GDPR, state privacy, SOC 2
Approved AI tools only	Central approval, access control, contract review, blocked/shadow AI process	HIPAA, CMMC, PCI DSS, SOC 2, ISO 27001

Control Area	Policy Evidence	Framework Alignment
Minimum necessary / minimization	Redaction, de-identification, tokenization, data owner approval	HIPAA, GDPR, state privacy, PCI DSS
BAA/DPA/AOC/vendor evidence	Legal/compliance vendor review before regulated data use	HIPAA, GDPR, PCI DSS, CMMC, GLBA
Human oversight	Output review, decision approval, override authority, documented rationale	NIST AI RMF, ISO 42001, EU AI Act, employment/consumer protection
Security controls	SSO/MFA, RBAC, encryption, logging, DLP, vulnerability management	NIST CSF, HIPAA Security Rule, CMMC, PCI DSS, SOC 2, ISO 27001
AI development controls	Threat model, testing, change management, model/update review, RAG access controls	NIST AI RMF, NIST CSF, ISO 42001, SOC 2
Incident response	AI incident categories, containment, evidence preservation, regulatory escalation	HIPAA, PCI DSS, CMMC, GDPR, state breach laws, SOC 2
Training	AI acceptable use, data handling, output verification, incident reporting	HIPAA workforce training, CMMC awareness, PCI security awareness, ISO 42001

Appendix G — 30/60/90-Day SMB Implementation Roadmap

Timeframe	Actions
0–30 days	Adopt policy; appoint Executive Sponsor and AI owner; identify approved/prohibited tools; launch AI inventory; block obvious unapproved tools for regulated data; publish workforce quick rules.
31–60 days	Complete vendor reviews for major AI tools; classify use cases; configure SSO/MFA/logging/retention; create intake workflow; train workforce; prioritize high-risk and regulated-data gaps.
61–90 days	Approve or deny pending use cases; complete compliance mapping; test incident playbook; conduct access review; launch metrics dashboard; schedule quarterly governance review.

Timeframe	Actions
Ongoing	Review inventory quarterly; review vendors annually; update for regulatory changes; monitor incidents; retest high-risk systems; refresh training annually.

Appendix H — Official Reference Sources

This policy was structured to align with recognized AI, cybersecurity, privacy, and compliance sources. Organizations should validate obligations against their specific contracts, laws, auditors, assessors, and counsel.

Source	Use in Policy	URL
NIST AI Risk Management Framework	AI risk governance, mapping, measuring, and managing AI risks	nist.gov/itl/ai-risk-management-framework
NIST AI 600-1: Generative AI Profile	Generative AI risk profile and companion to NIST AI RMF	nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence
ISO/IEC 42001:2023	Artificial Intelligence Management System requirements	iso.org/standard/42001
NIST Cybersecurity Framework 2.0	Cybersecurity risk governance and management outcomes	csrc.nist.gov/pubs/cswp/29/the-nist-cybersecurity-framework-csf-20/final
HHS HIPAA Security Rule	Administrative, physical, and technical safeguards for ePHI	hhs.gov/hipaa/for-professionals/security/index.html
32 CFR Part 170 — CMMC Program	CMMC program requirements for safeguarding FCI and CUI	ecfr.gov/current/title-32/subtitle-A/chapter-I/subchapter-G/part-170
DFARS CMMC rule	Contractual CMMC requirements and CMMC status/affirmation	federalregister.gov (2025-17359)
PCI SSC PCI DSS v4.0.1	Current PCI DSS version and effective-date guidance	blog.pcisecuritystandards.org/just-published-pci-dss-v4-0-1

Source	Use in Policy	URL
EU AI Act Service Desk Timeline	Progressive AI Act application timeline	ai-act-service-desk.ec.europa.eu/en/ai-act/timeline
FTC Artificial Intelligence resources	AI enforcement and consumer protection resources	ftc.gov/industry/technology/artificial-intelligence

Customization reminder

Before adoption, confirm the organization's AI scope, existing data classification, insurance requirements, contract flowdowns, customer commitments, regulatory scope, vendor list, incident reporting obligations, and approval authorities. Update placeholders and retain approval evidence.

— Template — customize for organization, regulatory scope, and counsel review —