

Enterprise Artificial Intelligence Use and Governance Policy

Industry Neutral Baseline for Responsible AI Use

Document control	Value
Organization	[Organization Name]
Policy owner	[Policy Owner or Function]
Approved by	[Board Executive or Governance Committee]
Effective date	[Month Day Year]
Review cycle	At least annually and upon material change
Classification	[Internal or Public]
Version	1.0

This policy establishes a common, risk-based foundation for the selection, development, acquisition, deployment and use of artificial intelligence. It is designed for adaptation by organizations of any size or industry. Sector-specific laws, contractual duties and professional obligations must be documented through the applicability process in this policy.

This document is a governance baseline, not legal advice and not a guarantee of compliance. Each organization must tailor it to its business, risk tolerance, technology environment and legal obligations.

Prepared as a reusable policy framework by

Peter Gallinari

Technology Cybersecurity Privacy and Governance Executive

1 Purpose

The purpose of this policy is to enable beneficial AI use while protecting people, information, operations, customers and the organization. It establishes accountable decision-making, proportionate controls and lifecycle oversight for AI systems and AI-enabled activities.

AI may improve productivity, analysis, accessibility, service delivery and innovation. It may also create or amplify privacy, security, legal, ethical, operational, safety, financial and reputational risks. The organization will therefore govern AI according to its actual capabilities, data access, intended use and potential impact—not merely its product name or vendor description.

2 Scope

This policy applies to all workforce members and relevant third parties, including employees, officers, board members, contractors, consultants, temporary personnel, interns, volunteers, service providers and vendors acting for or on behalf of the organization.

It applies to AI that is acquired, built, configured, integrated, tested or used for organizational purposes, including generative AI, machine learning, predictive systems, decision-support tools, embedded AI features, copilots, chatbots, autonomous or semi-autonomous agents, robotic process automation with AI capabilities and AI supplied by third parties.

- The policy applies whether the AI is free or paid, organization-managed or personally obtained, standalone or embedded within an existing platform.
- Personal use that has no connection to organizational information, systems, work or representation is outside scope, subject to other applicable policies.

3 Policy principles

Principle	Organizational expectation
Accountability	A named business owner remains accountable for the use case, data, decisions, outcomes and controls.
Lawful and appropriate use	AI must be used for an approved purpose consistent with applicable law, contracts, policy and professional obligations.
Risk proportionality	Controls, approval and monitoring must increase with the sensitivity, autonomy, scale and potential impact of the use case.
Data stewardship	Only authorized data may be used. AI inherits existing access and may expose weaknesses already present in the environment.
Human authority	Human involvement must be meaningful, informed and capable of changing, stopping or reversing an outcome when required.
Security and resilience	AI must be protected against misuse, manipulation, unauthorized access, insecure integration and operational failure.
Transparency	Material AI use and limitations must be disclosed to decision-makers and affected persons when appropriate or required.

Principle	Organizational expectation
Continuous governance	Approval is not permanent. Material changes require reassessment and, when applicable, reauthorization.

4 Definitions

Term	Definition
Artificial intelligence system	A machine-based system that, for explicit or implicit objectives, generates predictions, recommendations, decisions, content or other outputs that influence physical or virtual environments.
AI use case	A defined business purpose, workflow or decision in which AI is used.
Generative AI	AI that produces synthetic content such as text, code, images, audio, video or structured data.
AI agent	An AI-enabled system that can plan, select tools, access resources or take actions with some degree of autonomy.
Material change	A change to the model, purpose, tools, data, permissions, users, autonomy, scale, environment or consequences that may alter risk.
High-impact use	A use that may materially affect a person's rights, access, eligibility, employment, safety, finances, health, education or other significant interests.
System owner	The person accountable for the technical operation, integration, security and lifecycle management of an AI system.
Business owner	The person accountable for the business purpose, data use, decisions, outcomes and risk acceptance for an AI use case.

5 Governance and accountability

The organization shall establish an AI governance structure proportionate to its size and risk. Smaller organizations may assign multiple responsibilities to the same qualified individual, but accountability must remain explicit and conflicts of interest must be managed.

Role	Minimum responsibilities
Governing body or executive leadership	Approve risk appetite, policy and high-impact use; receive material risk and incident reporting; provide resources.
AI governance lead or committee	Maintain the governance process; review risk classifications and exceptions; coordinate legal, privacy, security, data, procurement and business review.

Role	Minimum responsibilities
Business owner	Define purpose and benefit; confirm authority; approve data and users; ensure human oversight; accept residual risk; monitor outcomes.
System owner or technology	Maintain architecture, configurations, integrations, access, testing, logging, continuity and technical documentation.
Data owner or steward	Authorize data use; confirm classification, quality, provenance, retention and sharing requirements.
Privacy legal compliance and security	Identify applicable obligations and risks; define controls; advise on notices, assessments, contracts, testing and incidents.
Workforce members	Use only approved tools and data; verify outputs; disclose AI use when required; report errors, incidents and unintended behavior.
Internal audit or independent assurance	Evaluate design and operating effectiveness according to risk and independence requirements.

- No AI system or use case may operate without a named business owner and, where applicable, a system owner and data owner.
- Business ownership may not be transferred to a vendor, model or technology function. Technology may operate the system; the business remains accountable for its use and outcomes.

6 AI inventory and registration

The organization shall maintain an inventory of AI systems and use cases. Discovery shall include formally procured systems, internally developed models, embedded AI capabilities, free online tools, pilots, proofs of concept, agents, automations and material vendor AI features.

Required inventory field	Examples
Identity and ownership	System name, use case, business owner, system owner, vendor, version
Purpose and users	Intended purpose, prohibited purpose, user groups, affected persons
Data and access	Inputs, outputs, data classification, repositories, permissions, retention, geographic processing
Capability and autonomy	Recommendations, decisions, content generation, tools, actions, ability to transact or communicate
Risk and approval	Risk tier, assessment date, approvers, conditions, residual risk, exception status
Lifecycle and monitoring	Deployment date, testing, metrics, incidents, material changes, review and retirement date

- Inventory records must be updated before deployment and whenever a material change occurs.

- The organization shall use reasonable discovery methods to identify unregistered AI. Self-reporting alone is not sufficient for higher-risk environments.

7 Risk classification

Every use case shall be classified before production use. The organization may rename or expand the tiers, but it shall not eliminate the requirement for risk-based differentiation.

Tier	Illustrative characteristics	Minimum disposition
Low	Administrative assistance using approved tools and non-sensitive data; no consequential decisions or external actions.	Business approval, approved-tool controls and user verification.
Moderate	Internal analysis, customer-facing content, sensitive data, integrations or outputs that may influence decisions.	Documented assessment, owner approval, privacy and security review as applicable, testing and monitoring.
High	Significant decisions, regulated data, safety implications, large-scale profiling, high autonomy, financial transactions or material external impact.	Formal impact assessment, multidisciplinary approval, enhanced testing, meaningful human authority, continuous monitoring and executive risk acceptance.
Prohibited	Unlawful, deceptive, discriminatory, unsafe, unauthorized or inconsistent with organizational values or obligations.	Do not acquire, develop, deploy or use.

8 Permitted use

AI may be used when the tool, purpose, data, users and controls have been approved at the appropriate risk level. Examples may include drafting, summarization, translation, accessibility support, coding assistance, research support, forecasting, classification and workflow assistance.

- Users must apply professional judgment, verify material facts and remain responsible for work product submitted or communicated in their name.
- AI-generated content must not be represented as independently verified, original human analysis when that representation would be misleading.

9 Prohibited use

- Use of unapproved AI with confidential, personal, regulated, privileged, proprietary or security-sensitive information.
- Use of AI to evade law, policy, contractual duties, records obligations, security controls, monitoring or approval requirements.
- Fully automated high-impact decisions unless expressly authorized by law, policy and the organization's high-risk approval process.
- Creation or distribution of deceptive impersonations, fabricated evidence, unlawful surveillance, malicious code, harassment or discriminatory content.

- Entering credentials, authentication secrets, encryption keys, private keys, access tokens or protected system configurations into an AI tool unless explicitly approved for that purpose.
- Using AI output as the sole basis for legal, clinical, employment, financial, safety or other consequential advice or action without qualified review and required authorization.

10 Data privacy and information governance

The organization shall determine what data an AI system can collect, access, infer, create, retain, disclose and use for training or service improvement. Data authorization shall be based on purpose, ownership, classification and applicable obligations—not merely on technical accessibility.

- Use the minimum data reasonably necessary for the approved purpose.
- Do not use live personal, confidential or regulated data in development, demonstrations, testing or non-production environments unless specifically authorized and protected to production-equivalent standards. Prefer synthetic, masked or de-identified data.
- Confirm whether prompts, files, outputs, telemetry and feedback are retained, reviewed by humans, shared with subprocessors or used to train models.
- Apply approved retention, deletion, legal hold, records management and data residency requirements to AI inputs and outputs.
- Assess privacy impacts when AI introduces new collection, inference, profiling, monitoring, biometric processing, automated decision-making or secondary use.
- AI access shall not be treated as proof that the underlying data permissions are appropriate. Excessive inherited access must be remediated.

11 Security requirements

- Use organization-approved identities, multifactor authentication, least privilege, segregation of duties and role-based or attribute-based access where appropriate.
- Restrict connectors, plug-ins, tools, APIs and data sources to those approved for the use case. Default-deny shall be considered for agentic and high-impact systems.
- Protect against prompt injection, data exfiltration, insecure output handling, model manipulation, malicious content, supply-chain compromise and unauthorized model or configuration changes.
- Log material prompts, actions, decisions, tool calls, approvals, overrides and administrative changes consistent with privacy and retention requirements.
- Establish incident response, backup, recovery, fail-safe operation and manual alternatives proportionate to business criticality.
- Security testing shall cover the complete system—including identities, data, integrations, tools and workflows—not only the underlying model.

12 Human oversight authority and accountability

Human presence is not automatically meaningful human control. Where human oversight is required, the designated person must have sufficient information, competence, time, independence and authority to

understand the proposed action, challenge the system, deny approval, stop execution and escalate concerns.

- Approval interfaces must present material inputs, assumptions, limitations, uncertainty and consequences in a form the reviewer can reasonably understand.
- The organization shall guard against automation bias, approval fatigue and workflows that make rejection impractical or merely ceremonial.
- Accountability for a consequential outcome may not be assigned to a person who lacked meaningful knowledge or authority over the action.
- Material changes invalidate prior assumptions and shall trigger reassessment and, when required, reauthorization.

13 Agentic AI and automated actions

AI agents create additional risk because they may select tools, sequence tasks, access changing information and act across systems. Approval of an initial objective does not constitute unlimited authorization for every later action.

- Define the agent's authorized purpose, data, tools, systems, transaction limits, communication channels, duration and escalation conditions.
- Use least-privilege service identities and short-lived credentials. Do not share human credentials with agents.
- Require human authorization before irreversible, external, high-value, safety-relevant or legally consequential actions unless a formally approved exception establishes equivalent safeguards.
- Apply rate, financial, data-volume and time limits; validate destinations and recipients; preserve auditable action histories; provide an immediate suspension mechanism.
- Test for unintended tool use, goal drift, recursive behavior, prompt injection, cross-system data leakage and failure to stop.

14 Development acquisition and third parties

AI risk shall be assessed before purchase, development or material renewal. Procurement convenience, pilot status or vendor reputation does not replace due diligence.

Review area	Minimum questions
Purpose and capability	What does the system do, what can it access, decide or execute, and what must it never do?
Data practices	What data is collected, retained, used for training, shared, transferred or available to vendor personnel and subprocessors?
Security and resilience	How are identity, encryption, isolation, logging, vulnerabilities, incidents, continuity and secure development addressed?
Model and quality	What are the model sources, limitations, evaluation methods, known failure modes, bias risks and change practices?

Review area	Minimum questions
Legal and contractual	What rights, restrictions, warranties, audit rights, indemnities, deletion duties, notices and regulatory requirements apply?
Exit and portability	Can the organization retrieve data, preserve records, transition services and terminate access without unacceptable disruption?

- Contracts shall address notification of material model, subprocessor, data-use, security and functionality changes.
- Vendor assurances must be validated according to risk. Marketing claims and generic certifications are not sufficient evidence for high-impact use.

15 Assessment approval and deployment

Before production use, the business owner shall document the use case, expected benefit, affected persons, data, system boundaries, failure consequences, legal and contractual requirements, risk classification, controls, testing, residual risk and monitoring plan.

Gate	Required evidence
1 Discover and register	Inventory record, named owners, purpose, users, data and system boundaries
2 Map obligations and impacts	Applicable legal, contractual, privacy, security, safety, accessibility and records requirements
3 Measure and test	Performance, limitations, misuse, security, privacy, bias, human factors and resilience testing proportionate to risk
4 Manage and approve	Controls, accountable approvers, residual risk acceptance, deployment conditions and rollback plan
5 Monitor and reauthorize	Metrics, logs, incidents, drift, complaints, vendor changes and material-change review

- A pilot or proof of concept is not exempt. Its data and environment must match the approved test conditions.
- The organization shall not move a use case to production solely because the technology performed well in a demonstration.

16 Testing and validation

- Define acceptance criteria before testing and use representative scenarios, including foreseeable misuse and edge cases.
- Evaluate accuracy, reliability, robustness, privacy, security, fairness, accessibility, explainability and operational impact as relevant to the use case.
- Test human-AI interaction, including whether users understand limitations, recognize uncertainty and can effectively intervene.

- Independently validate high-impact systems when the development team’s independence or evidence is insufficient.
- Document known limitations and ensure they are communicated to users and decision-makers.

17 Transparency communication and intellectual property

- Disclose material AI interaction or AI-generated content when required by law, contract, professional duty or organizational standards.
- Provide appropriate notices and channels for questions, corrections, complaints, appeals or human review when people may be materially affected.
- Respect copyright, trademark, confidentiality, license and attribution requirements. Do not assume AI output is accurate, unique, non-infringing or owned by the organization.
- Externally published AI-generated content must receive review proportionate to its subject, audience and potential harm.

18 Monitoring change management and reauthorization

AI systems and their environments change. Owners shall monitor performance, access, drift, incidents, complaints, security events, vendor changes and unexpected outcomes throughout the lifecycle.

Material-change trigger	Examples
Model or configuration	New model, major version, system prompt, tuning, safety setting or decision threshold
Tools and integrations	New connector, API, application, retrieval source, action or communication channel
Data and permissions	New dataset, broader population, higher classification, expanded retention or access rights
Purpose and population	New business purpose, geography, customer group, workforce use or affected person
Autonomy and consequences	Increased action authority, transaction value, scale, frequency, irreversibility or impact
Legal and vendor status	New obligation, contract term, subprocessor, hosting location, ownership or service condition

- Material changes shall be documented before implementation and routed for reassessment at the appropriate risk level.

19 Incident reporting and response

Workforce members shall promptly report suspected AI incidents, including unauthorized disclosure, harmful or discriminatory output, security compromise, unexpected autonomous action, material error, policy violation, intellectual-property concern or inability to stop the system.

- The incident response process shall address containment, evidence preservation, affected data and persons, legal notification, vendor coordination, corrective action and safe restoration.

- The organization shall support good-faith reporting and prohibit retaliation for raising an AI risk or refusing an unauthorized AI action.

20 Training and competence

- All covered persons shall receive role-appropriate training before using organizational AI and periodically thereafter.
- Training shall cover approved tools, prohibited data, verification, limitations, security, privacy, records, intellectual property, incident reporting and human accountability.
- Developers, reviewers, owners and high-impact decision-makers shall receive specialized training appropriate to their responsibilities.

21 Records evidence and audit

The organization shall retain sufficient evidence to demonstrate what was approved, how the system operated, what changed, who exercised authority and how material outcomes were reviewed. Records shall be proportionate to risk and maintained according to approved schedules and legal holds.

- High-impact records should include assessment results, testing, approvals, versions, data sources, material prompts or decision inputs, outputs, actions, overrides, incidents and monitoring results as applicable.
- Audit access shall be controlled. Logging must not create unnecessary secondary collections of sensitive information.

22 Exceptions

Exceptions must be documented, time-limited and approved by the policy owner and appropriate business, privacy, security, legal or executive authority based on risk. An exception request shall state the business need, scope, duration, affected data and systems, risks, compensating controls and accountable owner.

- Emergency use shall be reviewed promptly after the event and shall not become permanent through continued use.

23 Enforcement

Violations may result in suspension of AI access, removal of a system, corrective action, contract remedies or disciplinary action consistent with applicable law and organizational policy. The organization should distinguish good-faith error from reckless, deceptive or intentional misuse while addressing the underlying control failure in either case.

24 Policy review and maintenance

The policy owner shall review this policy at least annually and when material changes occur in the organization, its AI use, technology, risk environment or legal obligations. Policies, inventories, assessments and controls shall be updated together so that governance reflects actual operations.

25 Industry and legal applicability overlays

This core policy is intentionally cross-sectoral. Each adopting organization shall complete an applicability analysis and attach requirements that reflect its operations. The absence of a listed overlay does not mean an obligation is inapplicable.

Overlay category	Illustrative considerations
Healthcare and life sciences	Health information, clinical validation, patient safety, medical-device requirements, research and professional standards
Financial services and insurance	Consumer protection, fair lending, model risk, suitability, fraud, financial records and third-party oversight
Government and public sector	Administrative law, procurement, public records, civil rights, accessibility, security standards and due process
Employment and human resources	Hiring, promotion, performance, monitoring, accommodations, discrimination, notices and appeal rights
Education and children	Student records, age-appropriate design, parental rights, learning impacts and heightened protections for minors
Critical infrastructure and manufacturing	Safety, operational technology, availability, supply chain, physical impacts and sector regulation
Global or multistate operations	Jurisdictional scope, cross-border transfers, local notices, automated-decision rights and localization requirements

26 Supporting standards and guidance

The organization may align implementation with recognized frameworks and standards, including the NIST Artificial Intelligence Risk Management Framework and its Generative AI Profile, relevant NIST cybersecurity and privacy guidance, ISO IEC 42001, ISO IEC 23894 and other sector-appropriate standards. Framework alignment does not replace the organization's applicability analysis or legal advice.

As of the policy's initial publication, NIST describes the AI RMF as a voluntary, cross-sectoral framework organized around Govern, Map, Measure and Manage. Because standards and laws change, references must be verified during each policy review.

Appendix A Adoption checklist

- ☐ Replace all bracketed fields and assign accountable owners.
- ☐ Define the organization's AI risk tiers and approval authorities.
- ☐ Inventory known AI systems, embedded features, pilots, agents and unapproved use.
- ☐ Classify data and validate permissions before connecting AI.
- ☐ Document applicable laws, contracts, professional duties and sector overlays.
- ☐ Approve AI tools, data conditions and prohibited-use rules.

- ☐ Establish impact assessment, procurement, testing and change-control workflows.
- ☐ Define meaningful human authority for consequential decisions and agent actions.
- ☐ Enable logging, monitoring, incident response, suspension and retirement.
- ☐ Train the workforce and specialized roles.
- ☐ Set metrics and report material risk to leadership.
- ☐ Approve, publish, enforce and review the policy.

Appendix B Minimum AI use case record

Field	Organization entry
Use case and business purpose	
Business owner and system owner	
AI system vendor and model	
Users and affected persons	
Data sources classifications and owners	
Connections tools and actions	
Risk tier and rationale	
Legal privacy security and sector requirements	
Testing results and limitations	
Human authority and escalation	
Approvals conditions and residual risk	
Monitoring metrics and material-change triggers	
Incident contact and suspension procedure	
Review and retirement dates	

Appendix C Approval

Approval role	Name and title	Signature	Date
Policy owner			

Approval role	Name and title	Signature	Date
Executive approver			
Legal or compliance			
Privacy or security			

Recommended companion documents: AI Use Case Intake; AI Impact Assessment; Approved AI Tools Register; Vendor AI Due Diligence Questionnaire; AI Incident Procedure; AI Agent Authorization Standard; Data Classification and Handling Standard; and sector-specific legal overlays.