

ENTERPRISE

Artificial Intelligence Use Guidance

Practical, Industry-Neutral Guidance for Responsible AI Use

Document control	Value
Organization	[Organization Name]
Guidance owner	[Policy Owner or Function]
Related policy	Enterprise Artificial Intelligence Use and Governance Policy
Effective date	[Month Day Year]
Review cycle	At least annually and upon material change
Classification	[Internal or Public]
Version	1.0

How to use this guidance

This document explains how to apply the Enterprise AI Policy in everyday work. It does not replace the policy, law, contract requirements, professional judgment or organization-specific procedures. Where requirements differ, the stricter applicable requirement controls.

Prepared as a reusable guidance framework by
Peter Gallinari
Technology, Cybersecurity, Privacy and Governance Executive

1 Purpose and relationship to policy

The Enterprise AI Policy establishes mandatory authority, accountability, boundaries, approvals, exceptions and consequences. This guidance translates those requirements into practical steps for workforce members, managers, business owners, system owners, data stewards, developers, reviewers and third parties.

Policy → Guidance → Standards → Procedures

The policy states what is required. This guidance explains how to make responsible decisions. Standards define required control specifications. Procedures and workflows describe who performs each task, when and how. Technical controls operationalize and enforce those requirements.

Who should use this document

Audience	Use this guidance to
All workforce members	Choose approved tools, protect information, verify outputs and report concerns.
Managers and business owners	Define purpose, ownership, risk, approvals, human authority and monitoring.
Technology and developers	Design secure integrations, identities, testing, logs, limits and suspension controls.
Privacy, legal, compliance and security	Map obligations, assess impacts, define controls and advise decision-makers.
Procurement and vendor management	Evaluate AI capabilities, data practices, contracts, changes and exit requirements.
Audit and assurance	Evaluate whether governance and controls operate as approved.

Core message

AI inherits the organization's environment—including its data, permissions, gaps and assumptions. Responsible AI use begins by identifying what exists and understanding what the AI can access, infer, create, retain, disclose, decide or do.

You cannot manage what you cannot see.

Before governance can become continuous or machine-enforceable, the organization must identify the AI system or use case, its owners, data, repositories, permissions, integrations, users, actions and potential consequences.

2 Before using any AI tool

Do not begin with the prompt. Begin with the purpose, tool, data and authority.

Five questions to answer first

1. Is the AI tool or embedded feature approved for organizational use?
2. Is this specific business purpose approved, and is a named business owner accountable?
3. What information will the AI receive, retrieve, infer or generate—and how is it classified?
4. Could the output influence a person, decision, transaction, external communication or system action?
5. Does this use require registration, assessment, specialist review or additional authorization?

Do	Do not
Use the approved tool for the approved purpose.	Assume a familiar brand, paid license or embedded feature is automatically approved.
Use an organizational account and approved identity.	Use a personal account for organizational information or activity.
Confirm the tool's data conditions before entering information.	Assume prompts and files are private, temporary or excluded from training.
Register the use case when required.	Call an activity a pilot, experiment or proof of concept to avoid review.
Stop and ask when the classification, ownership or authorization is unclear.	Proceed because the information is technically accessible.

A simple stopping rule

If you cannot identify the approved tool, approved purpose, accountable owner and authorized data, do not use AI for the activity until those questions are resolved.

3 Discovering AI and determining what must be registered

AI may be visible as a chatbot or invisible inside software already in use. Registration should focus on actual capability and risk—not marketing labels.

Common places AI may appear

Location	Examples to look for
Productivity platforms	Copilots, writing assistants, meeting summaries, search, translation and content generation.
Business applications	CRM scoring, fraud detection, HR ranking, customer service, forecasting and recommendations.
Development environments	Code assistants, test generation, vulnerability analysis and autonomous development agents.
Data and analytics	Classification, profiling, prediction, natural-language query and synthetic data generation.
Security and operations	Alert triage, automated response, identity analytics, monitoring and workflow automation.
Free or personal tools	Public chatbots, browser extensions, mobile applications and consumer AI services.
Third parties	Vendor-managed models, embedded features, subcontractor tools and AI-enabled services.

Register the use case when

- AI accesses organizational data, systems, repositories, APIs or communications.
- AI creates external content, recommendations, decisions or actions on the organization's behalf.
- AI processes personal, confidential, regulated, privileged, proprietary or security-sensitive information.
- AI affects eligibility, employment, healthcare, education, safety, finance, legal rights or other significant interests.
- An AI agent or automation can select tools, contact people, alter records, initiate transactions or take system actions.
- A material change expands the model, purpose, users, data, permissions, tools, autonomy, scale or consequences.

Shadow AI

If you discover an unapproved tool or unknown AI feature, do not investigate by uploading sensitive information. Stop use when safe, preserve relevant facts and report it through [AI reporting channel]. Good-faith reporting should be supported.

4 Data discovery, classification and permission validation

An AI tool can surface, combine or infer information at a speed and scale that ordinary users cannot. Existing access does not prove that the access is appropriate for AI.

What to identify

Question	Evidence or action
What data is involved?	Identify prompts, uploads, repositories, retrieved content, outputs, logs, feedback and inferred data.
Where does it reside?	Identify applications, databases, cloud services, collaboration sites, file shares, endpoints, email and vendor environments.
Who owns it?	Obtain authorization from the accountable data owner or steward—not only the system administrator.
How is it classified?	Apply the organization’s classification and handling standard to inputs, outputs and derived information.
Who can access it?	Validate users, groups, service accounts, inherited permissions, links, external sharing and administrative access.
What happens to it?	Determine retention, training, human review, subprocessor access, residency, disclosure, deletion and legal-hold behavior.

Data-use rules

- Use the minimum data reasonably necessary for the approved purpose.
- Prefer synthetic, masked or de-identified data for development, demonstrations and testing.
- Do not place live personal, confidential or regulated data in non-production unless specifically authorized and protected to production-equivalent standards.
- Do not enter passwords, keys, tokens, authentication secrets or protected configurations unless the tool and purpose are explicitly approved.
- Treat AI-generated summaries, classifications and inferred attributes according to their content and impact—not as automatically non-sensitive.

Foundation sequence

Discovery → Classification → Privacy → Security → AI Governance → Enforcement. These activities may overlap, but enforcement cannot reliably cover systems or data the organization has not identified.

5 Safe and responsible prompting

A prompt is a disclosure and an instruction. Treat it with the same care as an email, file transfer, query or system command.

Good practice	Why it matters
State the task, audience, boundaries and required format.	Clear constraints reduce ambiguity and make review easier.
Remove unnecessary names, identifiers, secrets and sensitive details.	Data minimization limits exposure and secondary use.
Ask the AI to distinguish facts, assumptions and uncertainty.	Confident language is not evidence of accuracy.
Request sources when appropriate—and verify the sources yourself.	AI may invent or misstate citations, cases, policies or technical references.
Break consequential work into reviewable stages.	Review is more meaningful before an output becomes an action.
Preserve required records without over-logging sensitive content.	Auditability must be balanced with privacy and retention requirements.

Example: safer drafting request

Instead of

“Write a response to this customer complaint,” followed by the customer’s full record in an unapproved public tool.

Use

An approved tool, the minimum authorized facts, a clearly defined audience and tone, and an instruction to flag uncertainty. A qualified person must verify the response before it is sent.

Prompt injection and retrieved content

Documents, websites, messages and retrieved records may contain instructions designed to manipulate an AI system. Users and system designers should treat retrieved content as untrusted input, restrict tool permissions, validate destinations and require authorization before consequential actions.

6 Reviewing and verifying AI output

AI output is a draft, prediction or recommendation—not proof. The required level of review increases with sensitivity, impact, scale, autonomy and irreversibility.

Output type	Minimum review
Routine internal draft	Check accuracy, relevance, tone, confidentiality and consistency with organizational requirements.
Research or factual summary	Verify material facts, dates, quotations, sources and whether relevant information was omitted.
Code or technical configuration	Review, test and scan in an approved environment; validate dependencies, secrets, licensing and security behavior.
Customer-facing or public content	Confirm facts, authorization, disclosures, intellectual property, accessibility, privacy and potential harm.
Consequential advice or decision	Require qualified review, documented authority, appropriate testing and an appeal or escalation path when applicable.
Automated action	Confirm authorization boundaries, limits, destination, reversibility, monitoring, logs and immediate suspension capability.

Verification checklist

- Are the material facts independently supported?
- Are cited sources real, relevant and accurately represented?
- Is anything confidential, biased, harmful, misleading or unlawfully discriminatory?
- Did the AI omit important context, uncertainty or limitations?
- Would I be comfortable taking responsibility for this output under my name and role?
- Does another qualified person or approval authority need to review it?

7 Meaningful human authority

A person appearing in an approval chain does not automatically establish human control. The reviewer must understand the proposed action and retain practical authority to change, deny, stop or escalate it.

Human presence	Meaningful human authority
A person receives a notification.	The person receives understandable information about inputs, assumptions, uncertainty, limitations and consequences.
A button requests approval.	Rejecting or changing the action is practical and does not create unreasonable pressure or penalty.
A reviewer is assigned.	The reviewer has competence, time, independence and authority appropriate to the decision.
A human approved the original design.	Material changes to the model, tools, data, permissions, purpose or autonomy trigger reassessment and reauthorization.
Someone can be blamed afterward.	Accountability is assigned to people who had meaningful knowledge and authority before the outcome.

Before approving a consequential action

6. What exactly am I approving now?
7. Has the model, workflow, data or authority changed since the last review?
8. Can I see the material basis, uncertainty and likely consequences?
9. Can I deny, modify or stop the action before harm occurs?
10. Is escalation available—and do I have enough time to use it?

Guard against ceremonial approval

Design workflows to reduce automation bias and approval fatigue. High volume, poor explanations, rushed deadlines or difficult rejection paths can turn nominal oversight into a rubber stamp.

8 AI agents, copilots and automated actions

Agents require additional care because they may choose tools, retrieve changing information, sequence tasks, communicate externally and take actions across systems. Approval of a general objective is not unlimited authorization for every later action.

Define the authorization boundary

Boundary	Document
Purpose	Permitted objective, prohibited purposes and completion condition.
Data	Approved sources, classifications, repositories, volume and retention.
Tools and systems	Permitted connectors, APIs, applications, commands and destinations.
Actions	What the agent may recommend, draft, change, send, approve, purchase or execute.
Limits	Time, rate, financial value, data volume, geography, recipients and retry thresholds.
Human checkpoints	Actions requiring review or authorization before execution.
Stop and recovery	Immediate suspension, rollback, incident response and manual alternative.

Never rely on the agent to define its own authority

- Use least-privilege service identities and short-lived credentials.
- Do not share human credentials with an agent.
- Validate recipients and destinations before external communication or data transfer.
- Require human authorization before irreversible, high-value, safety-relevant or legally consequential actions unless formally approved safeguards provide equivalent control.
- Test for prompt injection, unintended tools, goal drift, recursive behavior, leakage and failure to stop.

9 Privacy, security, records and intellectual property

Privacy

Assess whether AI creates new collection, inference, profiling, monitoring, biometric processing, automated decision-making, secondary use or disclosure. Provide notices, choices, correction, appeal or human review when required.

Security

Use approved identities, multifactor authentication, least privilege and controlled integrations. Protect against prompt injection, exfiltration, insecure outputs, malicious content, supply-chain compromise and unauthorized configuration changes. Test the full system—not only the model.

Records

Apply approved retention, deletion, legal-hold and records-management rules to prompts, files, outputs, logs, decisions and actions. Preserve sufficient evidence for high-impact use without creating unnecessary secondary collections of sensitive data.

Intellectual property

Do not assume AI output is accurate, unique, non-infringing or owned by the organization. Respect copyright, trademark, license, confidentiality, attribution and contractual requirements. Avoid uploading third-party material unless authorized.

Situation	Guidance
Meeting transcription or summary	Confirm authorization, notice or consent requirements; limit participants, retention and downstream sharing.
Public marketing content	Verify claims, brand approval, rights, disclosures and whether people could be misled.
Source code assistance	Use approved repositories and tools; review generated code, licenses, dependencies and security.
Legal, clinical or financial material	Do not use AI output as the sole basis for advice or action; require qualified review.
Personal or regulated information	Confirm necessity, authorization, tool conditions, safeguards, retention and individual rights.

10 Third-party and vendor AI

A vendor's reputation, contract, certification or statement that its product is 'enterprise ready' does not establish that a particular use is appropriate. Evaluate the capability in the organization's actual environment.

Questions to ask

Area	Minimum questions
Capability	What can the system access, infer, generate, decide, communicate or execute? What must it never do?
Data	What is retained, used for training, reviewed by people, transferred, shared with subprocessors or available to administrators?
Security	How are identity, isolation, encryption, vulnerabilities, logging, incidents, continuity and secure development addressed?
Model quality	What are the sources, limitations, evaluation methods, known failure modes, bias risks and change practices?
Contract	What rights, restrictions, notices, audit rights, deletion duties, warranties, remedies and regulatory obligations apply?
Change	How will the organization learn of changes to models, subprocessors, hosting, data use, security or functionality?
Exit	Can data and records be retrieved, access terminated and services transitioned without unacceptable disruption?

Remember

Risk may be transferred contractually in part, but accountability for the organization's purpose, data, decisions and outcomes cannot simply be outsourced to the vendor.

11 Change, monitoring and reauthorization

AI approval is not permanent. A previously acceptable use may become higher risk when its model, data, integrations, users, purpose, autonomy or consequences change.

Change trigger	Examples
Model or configuration	New model or 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.
Purpose and population	New business purpose, geography, customer group, workforce use or affected person.
Autonomy and consequences	Greater authority, transaction value, scale, frequency, irreversibility or impact.
Legal or vendor status	New obligation, contract term, subprocessor, hosting location, ownership or service condition.

Monitor according to risk

- Performance, errors, drift and known limitations
- Access, permissions, data sources and tool calls
- Complaints, appeals, overrides and human-review effectiveness
- Security events, privacy issues, unexpected actions and inability to stop
- Vendor, model, contract, hosting and subprocessor changes
- Whether the use remains necessary, proportionate and aligned with its approved purpose

Reauthorization question

Does the current system still match what was originally assessed and approved? If not, document the material change and route it for reassessment before implementation.

12 Reporting concerns and AI incidents

Prompt reporting allows the organization to contain harm, preserve evidence, meet notification duties, correct controls and learn from failure. Report suspected issues even when all facts are not yet known.

Report immediately when AI

- Discloses or appears to expose information without authorization.
- Produces harmful, discriminatory, deceptive or materially incorrect output.
- Takes an unexpected action or acts outside its approved boundary.
- Cannot be stopped, suspended or reversed as expected.
- Uses an unapproved tool, data source, integration, account or credential.
- Shows signs of compromise, manipulation, prompt injection or data exfiltration.
- May have affected a person's rights, eligibility, employment, health, safety, finances or other significant interests.

What to do

Step	Action
1 Protect	Stop or suspend the activity when authorized and safe. Do not destroy evidence.
2 Report	Contact [AI incident channel / service desk / security operations]. For immediate safety risks, use [emergency contact].
3 Preserve	Record the tool, time, account, prompt or action, affected data, output and people involved without spreading sensitive information.
4 Cooperate	Support containment, privacy and legal review, vendor coordination, notification, corrective action and safe restoration.
5 Learn	Address the underlying process or control weakness—not only the individual action.

Good-faith reporting

People should be able to raise AI risks, question an output or refuse an unauthorized action without retaliation. Do not conceal mistakes or continue unsafe use to avoid embarrassment.

13 Practical scenarios

Scenario	Recommended response
An employee wants a public chatbot to summarize a confidential report.	Do not upload it. Use an approved tool and authorized data conditions, or remove protected content only if the resulting material is genuinely non-sensitive.
A licensed business platform activates an AI assistant by default.	Treat the feature as a new capability. Identify access, data use, retention, users, permissions and risk; register and approve it before organizational use.
A manager uses AI to rank job applicants.	Pause for high-impact review. Assess discrimination, accessibility, notices, data authority, testing, meaningful human authority, documentation and appeal requirements.
A developer accepts AI-generated code because it compiles.	Require code review, security testing, dependency and license checks, secrets scanning and validation in an approved environment.
An agent can email customers and issue refunds.	Define recipients, templates, value and rate limits, approval thresholds, logs and suspension. Require authorization before material or exceptional actions.
A reviewer sees an AI recommendation but lacks the evidence to challenge it.	Do not treat the review as meaningful oversight. Provide material inputs, assumptions, limitations and authority—or redesign the workflow.
A pilot succeeds with synthetic data and is ready for live use.	Production is a new risk decision. Reassess live data, permissions, scale, integration, affected persons, failure consequences, controls and monitoring.

14 Quick decision guide

Use these questions before proceeding. Any “No” or “Unknown” answer should be resolved at the appropriate risk level.

Decision question	Yes	No / Unknown
Is the tool approved for this purpose?	Continue.	Stop and seek approval.
Is a named business owner accountable?	Continue.	Assign ownership before use.
Are the data sources, classifications, owners and permissions known?	Continue.	Complete discovery and authorization.
Is the use registered and risk-classified when required?	Continue.	Register and assess it.
Are legal, privacy, security, records and contract duties addressed?	Continue.	Obtain specialist review.
Can material output be verified by a qualified person?	Continue.	Do not rely on the output.
For consequential action, is human authority meaningful?	Continue.	Redesign or prohibit the workflow.
For automation, are tools, actions and limits explicitly authorized?	Continue.	Restrict the system before use.
Can the system be monitored, suspended and safely recovered?	Continue.	Do not deploy until controls exist.

When in doubt

Contact [AI governance contact] or submit [AI use case intake link]. For suspected incidents, use [incident reporting channel].

Appendix A Workforce checklist

- ☐ I am using an organization-approved AI tool and account.
- ☐ My purpose is authorized and consistent with my role.
- ☐ I understand what information the AI will access, receive, infer, retain and disclose.
- ☐ The data is necessary, appropriately classified and authorized by its owner.
- ☐ I have not included passwords, tokens, secrets or unnecessary sensitive information.
- ☐ I understand the tool's important limitations and will verify material output.
- ☐ I will not use AI as the sole basis for consequential advice, decisions or actions.
- ☐ I will obtain required review before external publication or communication.
- ☐ I know how to report an error, unauthorized use, disclosure or unexpected action.
- ☐ I remain accountable for work submitted or communicated in my name.

Appendix B Manager and owner checklist

- ☐ Define the business purpose, expected benefit, users and affected persons.
- ☐ Assign business, system and data ownership.
- ☐ Register and risk-classify the use case before production.
- ☐ Identify data, repositories, permissions, integrations, actions and system boundaries.
- ☐ Map legal, privacy, security, records, safety, accessibility and contractual obligations.
- ☐ Define testing, acceptance criteria, human authority and escalation.
- ☐ Approve residual risk, deployment conditions, monitoring and rollback.
- ☐ Reassess material changes and retire access, data and records when use ends.

Appendix C Organization-specific information

Item	Organization entry
AI governance contact	[Name, team, email or portal]
AI use case intake	[Link or process]
Approved AI tools register	[Link or location]
Data classification standard	[Link or location]
Privacy / legal / compliance contact	[Contact]
Information security contact	[Contact]
AI incident reporting	[Channel, service desk or hotline]
Emergency suspension contact	[Contact]
Records and retention guidance	[Link or location]
Accessibility guidance	[Link or location]
Sector or jurisdiction overlays	[Links or attachments]

Appendix D Related documents

- Enterprise Artificial Intelligence Use and Governance Policy
- 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
- Sector-specific legal and regulatory overlays

Tailoring statement

This guidance is intentionally cross-sectoral. Each organization must adapt it to its size, services, risk tolerance, technology environment, workforce, contracts, professional duties and applicable laws. It is not legal advice or a guarantee of compliance.

Build the Foundation for AI, Privacy, Security, and Trust.