INTERNAL DOCUMENT TEMPLATE

Policy on the Use of Artificial Intelligence Systems

A safe path from idea to deployment

Field	Value to be completed
Organization	[ORGANIZATION NAME]
Document owner	[ROLE OR FUNCTION]
Approved by	[NAME OR BODY]
Version	[VERSION NUMBER]
Effective date	[DD.MM.YYYY]
Review date	[DD.MM.YYYY]
Document classification	[INTERNAL PUBLIC OTHER]

This template should be adapted to the organization's structure, business profile, information classification, applicable laws, and adopted risk management model. Before approval, the document should be reviewed by the appropriate legal, security, data protection, and risk management functions.

How to use this template

- Replace all fields in square brackets with the organization's details.
- Remove options that do not fit the adopted decision-making model.
- Complete the appendix listing approved tools and owners.
- Link this document to the information security, data protection, procurement, SSDLC, and incident management policies and procedures.

1. Policy purpose

The purpose of this policy is to enable employees and contractors to use artificial intelligence systems safely, lawfully, and in a controlled manner. The policy sets out rules for using AI tools, classifying use cases, assessing risk, making decisions, launching solutions, and responding to errors and incidents.

The organization does not limit AI governance to a list of prohibited products. It provides employees with a predictable route for submitting an idea, obtaining a rapid risk assessment, and receiving a decision on an experiment, pilot, or deployment.

2. Scope

This policy applies to [EMPLOYEES CONTRACTORS SUPPLIERS INTERNS OTHER PERSONS] using AI on behalf of the organization or using its data, devices, accounts, infrastructure, or processes.

In particular, this policy covers:

- public and corporate chatbots and AI assistants
- AI features embedded in office applications and SaaS services
- generation and analysis of text, images, audio, video, and code
- RAG solutions, semantic search engines, and corporate knowledge bases
- fine-tuned, local, and API-accessible models
- automations and AI agents that perform actions in systems
- applications and prototypes created through vibe coding or with coding assistants

3. Definitions

Term	Meaning used in this policy
AI system	A machine-based system operating with varying levels of autonomy that generates outputs such as predictions, content, recommendations, or decisions based on input data.
Generative AI	An AI system that generates new content, including text, code, images, audio, or video.
AI agent	A solution that plans steps, uses tools or data, and may perform actions that change the state of systems.
RAG	An architecture in which the model receives context retrieved from an external knowledge repository.
Shadow AI	Use of an AI system outside the approved process, without required registration, assessment, contract, controls, or approval.
Vibe coding	Creating code or applications with extensive AI assistance, often based on natural-language instructions.
Use case	A description of the specific purpose, process, data, recipients, and manner in which an AI system is used.
Use case owner	The person accountable for the business purpose, output quality, risk, and solution lifecycle.

4. Governing principles

- Purpose before tool. Every use case must have a justified purpose and an owner.
- Data minimization. Only data necessary to perform the task may be entered into a system.
- Output verification. AI-generated content and recommendations are not automatically treated as correct or binding.
- Least privilege. An agent or integration receives only the access required for the specific task.
- Human oversight. Decisions with a significant impact on people, finances, rights, or security are subject to appropriate human control.
- Accountability. The organization maintains an appropriate record of data, sources, model versions, actions, and approvals.
- Secure lifecycle. A prototype does not become a production solution until the specified requirements have been met.
- Report without concealment. The organization supports good-faith action and prompt reporting of errors or Shadow AI.

5. Data entered into AI systems

Whether data may be entered depends on its classification, the supplier agreement, processing location, retention, use for model training, access controls, and the approved use case.

Data class	Default rule	Conditions or examples
Public	Permitted in approved tools	Content officially published or intended for publication. Copyright, reliability, and purpose must still be checked.
Internal	Conditionally permitted	Only in tools approved for this data class, using a business account and appropriate privacy settings.
Confidential	Requires submission and approval	Contracts, commercial information, code, architecture, customer data, and non-public strategies. Assessment by the data owner and Security is required.
Personal data	Requires a legal basis and privacy assessment	The purpose, legal basis, minimization, retention, roles of the parties, and any need for a DPIA must be determined.
Special-category data	Prohibited by default without formal approval	Permission requires an assessment by the DPO or Legal, the data owner, and an approved environment.
Secrets and credentials	Prohibited	Passwords, tokens, API keys, private keys, recovery codes, and data enabling takeover of an account or service.
Data subject to secrecy or sector restrictions	Prohibited without a specific decision	Regulated data, professional secrets, classified information, or contractually restricted data.

If a user cannot determine the data class or supplier terms, they must not enter the data into the system and must contact [DATA OWNER SECURITY DPO SERVICE DESK].

6. Approved tools

The current list of tools is provided in Appendix 1 or in [REGISTER NAME AND LOCATION]. Inclusion of a product in the list does not constitute approval for every use. Each approval specifies the permitted accounts, data classes, features, integrations, region, retention, and restrictions.

Status	Meaning	User action
Approved	The tool may be used within the described scope.	Follow the entry conditions and submit a new use if it meets the criteria in Section 8.
Conditional	Permitted only for the specified team, data, pilot, or period.	Confirm that your case falls within the approved scope.
Pending assessment	The assessment or contract has not been completed.	Do not use organizational data. Submit a request.
Prohibited	The tool or feature may not be used in the organization.	Do not use it. If the tool is already in use, report this in accordance with Section 14.

7. Prohibited uses

AI must not be used for a purpose or in a manner that:

- violates the law, individual rights, contractual obligations, sector regulations, or the organization's internal rules
- carries out practices prohibited under the applicable provisions of the AI Act or other regulations
- requires disclosure of secrets or credentials, or bypassing safeguards
- enables covert profiling, discrimination, manipulation, or monitoring of individuals without an appropriate legal basis and authorization
- generates or distributes content impersonating individuals or the organization without authorization and required labeling
- automatically makes or executes decisions with a significant impact without approved human oversight
- allows an agent to make payments, issue refunds, arrange shipments, delete data, change permissions, or perform other high-impact actions without service-side controls
- combines access to confidential data with unrestricted internet browsing or transmission of data to external addresses
- deploys a vibe-coded application to production outside the approved SSDLC
- uses data or materials without the required license, legal basis, or authorization

8. When a use case must be submitted

Submission is required before a pilot or use begins if at least one of the following applies:

- a new tool, model, supplier, account, plug-in, or extension
- non-public data or personal data
- RAG, a vector index, model fine-tuning, or a proprietary knowledge base
- integration with email, files, CRM, ERP, HR, code, a database, or an API
- an agent that performs actions or makes multi-step decisions
- an impact on employment, education, access to services, credit, health, safety, or individual rights

- automated generation of external communications, offers, contracts, production code, or decisions
- use of biometrics, emotion recognition, children's data, or data concerning particularly vulnerable groups
- a material change to the purpose, data, model, permissions, integration, autonomy, or scale of an existing solution

Submission is not required for uses listed in the approved low-risk catalog, provided that the user meets all catalog conditions. The catalog and exceptions are maintained by [REGISTER OWNER].

9. A safe path from idea to deployment

Each submitted use case follows a path proportionate to its risk. The purpose of the process is to provide a prompt response: approval, conditional approval, a requirement for additional testing, or a reasoned refusal.

Stage	Required outcome	Exit criterion
1 Idea	Purpose, owner, users, expected outcome, and impact of error.	Complete submission using Appendix 2.
2 Triage	Initial risk threshold and list of involved roles.	Defined path and decision deadline.
3 Classification	The organization's role, AI Act classification, privacy, data, security, and business risk.	Documented requirements and risk owner.
4 PoC	Test in a limited environment with success and stop criteria.	No unacceptable results and known limitations.
5 Safeguards	Technical, organizational, legal, and process controls.	Pre-launch checklist completed.
6 Decision	Approval, conditional approval, stop, or prohibition.	Decision by a duly authorized person.
7 Production	Controlled launch, monitoring, instructions, rollback, and review date.	Acceptance and registration.
8 Maintenance	Monitoring of quality, cost, incidents, changes, and compliance.	Periodic reassessment or retirement.

10. Risk thresholds and decision times

Threshold	Example characteristics	Decision and indicative timeframe
Low	Public or ordinary internal data, no actions in systems, and output verified by a person.	Automatic approval from the catalog or by the process owner. [1 BUSINESS DAY].
Elevated	Confidential data, RAG, business recommendations, read-only integrations, or limited impact.	Business and data owners after consultation with IT, Security, the DPO, or Legal. [5 BUSINESS DAYS].
High	Sensitive data, decisions affecting people, write operations, financial actions, high autonomy, or an impact that is difficult to reverse.	A committee or risk-accepting person after formal assessment and testing. [15 BUSINESS DAYS].
Unacceptable	A practice prohibited by law or a risk that cannot be reduced to an acceptable level.	Prohibition or stop. Escalation to [DECISION-MAKING BODY].

Values and timeframes should be adapted to the organization's scale. Each threshold must identify one role with final decision-making authority. Failure to respond by the deadline does not constitute automatic approval.

11. Roles and responsibilities

Role	Decides or is accountable for	Does not delegate
Management board or sponsor	Risk appetite, resources, and oversight of the AI management system.	Accountability for establishing effective rules.
Use case owner	Purpose, success metric, output quality, users, and impact of error.	Accountability for the business outcome.
Data owner	Classification, access, minimization, retention, and permitted use of data.	Decisions on the scope of data.
IT or Product	Architecture, integrations, maintenance, versions, monitoring, and retirement.	Operational ownership of the solution.
Security	Threat model, required controls, testing, and cybersecurity risk.	Acceptance of business risk.
DPO or Legal	Privacy, legal basis, contracts, the AI Act, and other obligations.	Technical decisions outside its mandate.
Procurement or Vendor Management	Supplier assessment, contract terms, subcontractors, region, and service termination.	Assessment of the actual use case.
Risk Owner	Acceptance of residual risk within the granted mandate.	A decision without documented conditions and a review date.
User	Use within scope, verification of outputs, and reporting of issues.	Responsibility for informed use of the tool.

12. Pre-launch requirements

The scope of requirements is proportionate to risk, but before launch the following must be documented at a minimum:

☐ use case owner, data owner, and maintenance owner
☐ purpose, users, inputs, outputs, and prohibited behaviors
☐ the organization's role and the outcome of the regulatory classification
☐ assessment of the supplier, contract, retention, region, subcontractors, and use of data for training
☐ architecture, trust boundaries, integrations, permissions, and data flows
☐ a threat model and tests covering erroneous, malicious, and unexpected inputs
☐ a source verification mechanism and measures to limit hallucinations
☐ human oversight appropriate to the consequences of a decision or action
☐ logging of model versions, sources, tool calls, decisions, and approvals
☐ limits on cost, number of steps, runtime, and operation frequency
☐ monitoring of quality, security, cost, model changes, and incidents
☐ emergency shutdown, rollback, recovery, and retirement procedures
☐ user instructions, required training, and reassessment criteria

13. Additional requirements for AI agents

- Content from users, documents, websites, and messages is treated as untrusted data, not as system instructions.
- Each tool and service independently verifies identity, scope of access, object ownership, and authorization for the operation.
- Permissions are limited to the specific task, record, time period, and environment.
- High-impact operations require a preview and informed approval by an authorized person.
- Internet access is separated from secrets and confidential data; destinations and output parameters are controlled.
- The agent has step, cost, and time limits, as well as a technical kill switch.
- Logs make it possible to reconstruct the input, retrieved sources, decisions, tool calls, output, and approving person.

14. Reporting errors and incidents

Everyone must promptly report suspected data disclosure, an erroneous decision, an unauthorized operation, output manipulation, discrimination, incorrect content labeling, excessive cost, loss of control over an agent, or use of AI outside the approved scope.

Channel	Details to be completed
Form or ticketing system	[ADDRESS OR SYSTEM NAME]
Email address	[TEAM ADDRESS]
Phone for urgent incidents	[PHONE NUMBER]
Process owner	[ROLE OR TEAM]
Target response time	[TIME BY SEVERITY]

Where safe, the reporting person should stop the activity and preserve the conversation identifier, time, model, prompt, attachments, output, and actions performed. Evidence must not be deleted, and the person suspected of an attack must not be contacted independently.

Good-faith reporting, including disclosure of unapproved AI use, must not be discouraged. The priority is to limit harm, understand actual use, and safely regularize the process, without prejudice to accountability for deliberate violations.

15. Reassessment and retirement

Each register entry includes the date of the next assessment. The default frequency is [12 MONTHS FOR LOW RISK], [6 MONTHS FOR ELEVATED RISK], and [3 MONTHS FOR HIGH RISK], or more frequently if required by law, contract, or the risk owner's decision.

Reassessment is also required after:

- a change to the model, supplier, service terms, region, or retention rules
- a change to the purpose, user group, scale, data, integrations, tools, or autonomy
- a significant incident, error, complaint, audit finding, or indication of discrimination
- a new vulnerability, change in the threat profile, or change to safeguards

- a change in law, guidance, AI Act classification, or the organization's role
- a decline in quality, increase in cost, lack of an owner, or loss of the ability to maintain the solution

A solution is retired when it has no owner, fails to meet the conditions of the decision, its risk cannot be effectively monitored, it no longer fulfills its purpose, or its risk exceeds the organization's accepted appetite.

16. Exceptions

An exception to this policy requires a written request stating the purpose, scope, duration, owner, risk, compensating controls, and exit plan. The exception is approved by [ROLE OR BODY] after consultation with the relevant functions. An exception may not permit a practice prohibited by law.

17. Competence and AI literacy

The organization provides training appropriate to the role, knowledge, manner of AI use, and the system's impact on individuals. The minimum curriculum covers data protection, output verification, model limitations, prompt injection, safe use of tools, incident reporting, and the rules in this policy.

Use case owners, developers, administrators, persons approving outputs, and members of decision-making bodies receive advanced training tailored to their responsibilities.

18. Compliance with the AI Act and other requirements

For each submitted solution, the organization determines whether it meets the definition of an AI system, the role performed by the organization, the intended use case, and the applicable legal obligations. The analysis covers, in particular, prohibited practices, high-risk systems, transparency requirements, general-purpose AI models, deployer obligations, and any required fundamental rights impact assessment.

AI Act classification does not replace cybersecurity, privacy, and business risk assessments. A system outside the high-risk category may still cause significant harm to the organization or individuals.

The legal position and the timetable for application of the rules must be verified on the adoption date and during each policy review. As of 10 September 2026, the AI Act applies in phases, and the timetable reflects changes introduced by the AI Omnibus.

19. Enforcement and accountability

A violation of this policy may result in access restrictions, retirement of the solution, corrective action, or consequences provided for under internal rules and the law. The assessment takes into account the nature of the violation, intent, impact, cooperation in limiting harm, and promptness of reporting.

20. Policy approval

Role	Name	Decision	Date
Policy owner	[COMPLETE]	[APPROVAL]	[DATE]
Security	[COMPLETE]	[APPROVAL]	[DATE]
DPO or Legal	[COMPLETE]	[APPROVAL]	[DATE]
Management board or approving body	[COMPLETE]	[APPROVAL]	[DATE]

Appendix 1. Register of approved tools

Tool and owner	Status and scope	Data and accounts	Controls and review
[NAME VERSION SUPPLIER] Owner [ROLE]	[APPROVED CONDITIONAL PROHIBITED] Features [SCOPE]	Data classes [SCOPE] Accounts [BUSINESS OTHER]	Retention [PERIOD] Training [YES NO] Review [DATE]
[COMPLETE]	[COMPLETE]	[COMPLETE]	[COMPLETE]
[COMPLETE]	[COMPLETE]	[COMPLETE]	[COMPLETE]

Appendix 2. Use case submission form

Use case name

.....

.....

Business owner and function

.....

.....

Problem and expected outcome

.....

.....

Users and persons affected by the output

.....

.....

Input data and classification

.....

.....

System output and how it will be used

.....

.....

Tool model supplier and version

.....

.....

Integrations permissions and possible actions

.....

.....

Does the system use RAG memory internet access or tools

.....

.....

Potential impact of error or misuse

.....

.....

Proposed human oversight

.....

.....

Planned pilot date and scale

.....

.....

Appendix 3. AI risk assessment card

Area	Control questions	Outcome and evidence
Regulation	Is it an AI system? What is the organization's role? What is the use case? Does it involve a prohibited practice, high risk, transparency, GPAI, or FRIA?	[COMPLETE]
Rights and privacy	Who is affected by the output? What data is processed? Is a DPIA required? Is there a risk of discrimination or restriction of rights?	[COMPLETE]
Cybersecurity	What are the trust boundaries, data, permissions, integrations, vulnerabilities, misuse scenarios, monitoring, and response?	[COMPLETE]
Business	What are the impact of error, extent of harm, reversibility, cost, risk owner, and pilot conditions?	[COMPLETE]

Rating scale

Rating	Likelihood	Impact
1	Rare or highly unlikely	Minor and easily reversible
2	Unlikely	Limited
3	Possible	Significant
4	Likely	Severe
5	Very likely or expected	Critical or difficult to reverse

Risk register

ID	Scenario and asset	P	W	Score	Controls owner and deadline
R1	[THREAT LEADS TO IMPACT]	[1 5]	[1 5]	[L X I]	[CONTROLS ROLE DATE]
R2	[COMPLETE]				
R3	[COMPLETE]				
R4	[COMPLETE]				

Decision	Content
Initial rating	[LOW ELEVATED HIGH UNACCEPTABLE]
Residual risk	[RATING AFTER CONTROLS]
Decision	[APPROVAL CONDITIONAL APPROVAL STOP PROHIBITION]
Conditions and restrictions	[SCOPE ENVIRONMENT USERS DATA TIME]
Risk owner	[NAME ROLE]
Reassessment date	[DATE OR EVENT]

Appendix 4. Pre-launch checklist

☐ The use case has business, data, and maintenance owners.
☐ The purpose, users, data, sources, and expected output are described.
☐ The organization's role and AI Act classification have been recorded.
☐ The privacy, contract, and supplier assessments have been completed.
☐ The architecture, trust boundaries, and data flows are documented.
☐ Permissions correspond to specific tasks and are technically enforced.
☐ Functional, security, and AI-specific tests have been conducted.
☐ Prompt injection, leakage, poisoning, hallucinations, and tool misuse have been tested.
☐ Human oversight and actions requiring approval have been established.
☐ Logs allow reconstruction of sources, calls, output, and decision.
☐ Cost, step, time, and frequency limits are operational.
☐ Monitoring, an incident procedure, a kill switch, and rollback are in place.
☐ Users have received instructions and training.
☐ Residual risk has been accepted by the appropriate person.
☐ The scope, conditions, and reassessment date have been recorded in the register.

Appendix 5. Error or incident report form

Field	Information
Date and reporting person	[COMPLETE]
Tool system model and version	[COMPLETE]
Conversation or operation identifier	[COMPLETE]
Description of expected and actual behavior	[COMPLETE]

Field	Information
Data sources prompt and attachments	[COMPLETE OR INDICATE EVIDENCE LOCATION]
Actions performed by the system	[COMPLETE]
Persons systems and data affected by the incident	[COMPLETE]
Immediate actions taken	[COMPLETE]
Was the solution stopped	[YES NO NOT APPLICABLE]
Response owner and status	[COMPLETE]

Appendix 6. Sources and related documents

- AI Act Service Desk AI Act implementation timeline
- AI Act Service Desk Article 4 AI literacy
- Regulation EU 2024 1689 on EUR Lex
- NIST AI Risk Management Framework Generative AI Profile
- OWASP Top 10 for LLM and Generative AI Applications
- OWASP AI Agent Security Cheat Sheet
- NIST Secure Software Development Framework

Organizational documents to be linked

- [INFORMATION SECURITY POLICY]
- [DATA CLASSIFICATION AND PROTECTION POLICY]
- [INCIDENT MANAGEMENT PROCEDURE]
- [SSDLC AND CHANGE MANAGEMENT STANDARD]
- [PROCUREMENT AND SUPPLIER ASSESSMENT POLICY]
- [PERSONAL DATA PROTECTION POLICY]
- [REGISTER OF AI SYSTEMS AND USE CASES]