Confidential

Generative AI Use Policy

Version	1.0
Effective date	Month D, YYYY
Next review	Next review date
Scope	The entire Company (including group companies and contractors)
Owner	IT Department / Digital Transformation Office
AI stance	Balanced

Prepared	Reviewed	Approved
IT Department / Digital Transformation Office	AI Governance Committee	Executive Officer in charge of IT

This edition follows EU law. It was prepared with reference to the EU AI Act (Regulation (EU) 2024/1689, as amended by Regulation (EU) 2026/1744), the General Data Protection Regulation, the guidance of the European Commission and the European Data Protection Board, and other EU rules, as checked on October 30, 2026.

How to use this template

This document is a generative AI use policy template that any company can use as a starting point. Follow the three steps below to turn it into your company's own document. Delete this page when you are done.

1. **Fill in the colored fields.** Enter your company's details, such as the company name, dates, department names and contact point, in the places marked in this color.
2. **Follow the editor's notes.** The editor's notes in the text tell you what to replace. Delete each note once you have applied it.
3. **Remove chapters and lines that do not apply.** You may delete entire chapters, such as the industry-specific considerations, overseas regulations, procurement and cost chapters. If you remove a chapter in the builder (Shishin Kobo), the table of contents, chapter numbers and cross-references between chapters are renumbered automatically. If you delete chapters by hand, renumber them yourself.

> **Laws and figures**
>
> The EU laws in Chapter 1 and figures such as log retention periods and reporting deadlines are as of the time this document was prepared. National laws of the Member States, such as employment law and the rights of works councils, are not covered in detail. Before distribution, check them against your company's own rules and the latest laws, and adjust them as needed.

Contents

1. Introduction
2. Generative AI Basics
3. Scope and Audience
4. Definitions
5. Basic Principles
6. Usage Rules
7. Data Handling Rules
8. Quality Control of AI Output
9. Copyright and Intellectual Property
10. Rules for AI Agents
11. Security
12. Privacy Protection
13. Risk Management
14. Industry-Specific Considerations
15. Regulations Outside the EU and Global Operations
16. Procurement, Contracts and Vendor Management
17. Developing Systems with Built-in AI
18. Organization and Governance
19. Education and AI Literacy
20. Incident Response
21. Consequences of Violations
22. Cost and Environmental Considerations
23. Periodic Review and Revision Cycle
24. Laws and Guidelines Referenced
25. Frequently Asked Questions (FAQ)
26. Appendix: Fill-in Templates

1. Introduction

Revision history

Editor's note: When distributing a revised edition, answer the question "Date of first issue" and the table splits into a row for the first edition and a row for the revision. Describe the changes in your own words.

Version	Date	Main changes
Version 1.0	Effective date	First edition

1.1 Purpose of this Policy

Generative AI can greatly increase productivity, but it also carries risks such as leaks of confidential information, copyright infringement, hallucinations (output that is not factual) and discriminatory output. Since 2025, AI agents and physical AI have moved into practical use, and AI increasingly makes judgments and takes actions on its own.

This Policy sets out the standards of conduct that allow all employees of Company name and its affiliates to use generative AI safely and effectively. The aim is to encourage active use within clear rules, rather than to discourage use through prohibitions.

> **AI stance of the Company (Balanced)**
>
> Measures are matched to the level of risk. Low-risk use is free, limited-risk use requires approval, and high-risk use is strictly controlled, balancing use and control.

1.2 Status and legal basis

This Policy has been prepared in accordance with the following EU laws and guidance. Where national law of a Member State in which The Company operates sets stricter rules, national law applies in addition.

Law or guidance	Issued by	Status
Artificial Intelligence Act (Regulation (EU) 2024/1689)	European Parliament and Council	In force since August 1, 2024 and applied in stages. Prohibited practices and AI literacy apply from February 2, 2025, transparency obligations from August 2, 2026, and high-risk obligations from December 2, 2027 (see Chapter 13)
Digital Omnibus on AI (Regulation (EU) 2026/1744)	European Parliament and Council	In force since July 27, 2026. Postponed the high-risk obligations, revised the AI literacy duty in Article 4, and added two prohibited practices from December 2, 2026

General Data Protection Regulation (GDPR, Regulation (EU) 2016/679)	European Parliament and Council	Applies to any processing of personal data in AI tools, including prompts, uploaded files and outputs about individuals
Guidelines on prohibited AI practices and on the definition of an AI system	European Commission	Published in February 2025
Guidelines on the transparency obligations of Article 50	European Commission	Published on July 20, 2026, with a voluntary Code of Practice on transparency of AI-generated content (June 10, 2026)
Directive on copyright in the Digital Single Market (Directive (EU) 2019/790)	European Parliament and Council	Articles 3 and 4 set the text and data mining exceptions and the right of rights holders to opt out
Trade Secrets Directive (Directive (EU) 2016/943)	European Parliament and Council	Information stays a trade secret only while reasonable steps are taken to keep it secret. Careless entry into AI tools can undermine this
ISO/IEC 42001:2023	ISO/IEC	International standard for AI management systems

Editor's note: If **The Company** is a financial entity, an essential or important entity under NIS2, or a public body, add the sector rules that apply to you to this table. The points to note are in Chapter 14.

1.3 Revision policy

AI technology and regulation change quickly. As a rule, this Policy is revised every six months, with ad hoc revisions when major legal amendments or technological shifts occur. Revisions are recorded and notified to all employees.

2. Generative AI Basics

2.1 How generative AI works (key points)

Large language models (LLMs) learn from vast amounts of text the probability of which word comes next. They do not understand the meaning of a question; they statistically generate a plausible continuation. That is why they can be fluent and still wrong (hallucination). Understanding this property is the starting point for safe use.

2.2 Types and what they can do

Type	Main uses	Examples
Text generation	Writing, summarizing, translating, research support	ChatGPT, Claude, Gemini
Code generation	Implementation, review and debugging support	GitHub Copilot, Cursor, Claude Code
Image, video and audio generation	Creating assets, design support	Firefly, DALL-E, various video and audio AI tools
AI agents	Autonomous execution of multi-step tasks, business process automation	Various agent platforms

2.3 Limitations (do not over-rely on AI)

- AI does not know recent information or facts specific to the Company (it can answer only within what it was trained on and what it is given).
- It is prone to errors in calculations, strict fact-checking and citing sources. Always verify figures, proper nouns and URLs.
- When asked for its "reasoning," it does not reveal its internal decision process; it merely produces a plausible explanation after the fact.

3. Scope and Audience

Category	Covered	Level of application
Full-time and contract employees	All departments and all positions	All provisions apply
Temporary staff	Persons engaged in work for the Company	All provisions apply (the staffing agency is also notified)
Contractors	Outsourced work that handles data of the Company	Data and security provisions apply as contract terms
Group companies	Consolidated subsidiaries and affiliates	Each company prepares its own version based on this Policy
Officers	Directors, executive officers and auditors	All provisions apply (including the responsibilities in the governance chapter)

Editor's note: Adjust the level of application for contractors to match your contracting practice. Whether you impose all provisions or only the data and security provisions changes how the contract is worded.

3.1 AI covered by this Policy

This Policy covers anything that falls under any of the following, regardless of how it is provided. AI is defined by function, because listing tool names would miss new services.

Type	Covered
Text generation	Conversational AI services and features that write, summarize or translate text, including similar features built into business systems
Image, audio and video generation	Image generation, speech synthesis, video generation, and generative editing of existing material
Code generation	Completion and generation features built into development environments, and code review support
AI agents	Forms in which AI makes its own decisions and operates external systems (see Chapter 10)
AI developed in-house	AI systems built by the Company, including systems that let AI refer to internal documents (RAG)
Features added to existing products	AI features later added to business software used every day. They become subject to this Policy as soon as they are added

> **Operate on the assumption that the scope will keep growing**
>
> AI features are being added to familiar business software, and people increasingly start using them without realizing it. Not only new tools, but also new features added to existing tools are subject to the review in Chapter 6.

4. Definitions

Where the EU AI Act or the GDPR defines a term, this Policy uses that definition.

Term	Definition
AI system	A machine-based system that operates with some autonomy, may adapt after deployment, and infers from its input how to generate outputs such as predictions, content, recommendations or decisions (Article 3(1) of the EU AI Act).
Generative AI	AI systems that generate text, images, audio, video, code and other content. Most are built on general-purpose AI models.
General-purpose AI model	An AI model trained on large amounts of data that can perform a wide range of tasks and can be integrated into many systems (Article 3(63)).
Provider	A person or organization that develops an AI system or model, or has it developed, and places it on the market or puts it into service under its own name (Article 3(3)).
Deployer	A person or organization that uses an AI system under its own authority in a professional activity (Article 3(4)). When employees use approved AI tools for work, the Company is the deployer.
High-risk AI system	An AI system used in the areas listed in Annex III of the EU AI Act, such as employment, education and credit scoring, or as a safety component of products covered by Annex I.
AI agent	An AI system that perceives its environment and acts autonomously to achieve a goal. It coordinates, analyzes and executes tasks in cooperation with multiple systems.
Personal data	Any information relating to an identified or identifiable natural person (Article 4(1) of the GDPR). Names in a prompt, a customer email pasted for summarizing, and a photo of a person are all personal data.
Special categories of personal data	Data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs or trade union membership, and genetic data, biometric data used for identification, health data and data concerning sex life or sexual orientation (Article 9 of the GDPR).
Pseudonymization	Processing personal data so that it can no longer be attributed to a person without additional information that is kept separately (Article 4(5) of the GDPR). Pseudonymized data is still personal data.
Prompt	The instruction given to generative AI. A distinction is made between system prompts and user prompts.

Hallucination	A phenomenon in which AI plausibly outputs information that is not factual.
RAG (retrieval-augmented generation)	A technique that retrieves relevant information from external data and passes it to an LLM to improve the accuracy of answers.
Human oversight	Arrangements that let competent people understand, monitor and, where needed, override or stop the output of an AI system. Also called Human-in-the-Loop.
Prompt injection	An attack that uses malicious input to bypass controls and cause unintended behavior. Includes indirect attacks through documents and web pages.

5. Basic Principles

5.1 Seven requirements for trustworthy AI

The seven requirements in the Ethics Guidelines for Trustworthy AI, which informed the EU AI Act, are interpreted and applied in the context of the Company's business.

No.	Requirement	Interpretation and application
1	Human agency and oversight	AI output is reference information, and people make the final decisions. AI is not used to manipulate decisions or exploit emotions. People who oversee AI have the competence and authority to override it.
2	Technical robustness and safety	Tools are approved before use, tested for the task, and stopped when they behave unexpectedly. Models and their providers are checked before adoption.
3	Privacy and data governance	Personal data is entered into AI only with a legal basis, only as much as needed, and preferably pseudonymized. Confidential data stays in tools approved for it.
4	Transparency	People are told when they are interacting with AI, and AI-generated content is labeled where EU law or this Policy requires it. Decisions supported by AI can be explained.
5	Diversity, non-discrimination and fairness	Output is checked so that discriminatory results do not influence decisions about people. Tools are accessible to staff with disabilities.
6	Societal and environmental well-being	AI use takes account of its effect on employees, customers and society, and of its energy use.
7	Accountability	Responsibilities for each AI use are assigned, usage logs are kept, and concerns can be raised without fear of reprisal.

5.2 Risk-based approach

Like the EU AI Act, the Company does not regulate all AI use uniformly. It sorts each use into one of four levels and sets measures accordingly. Chapter 13 explains how to classify a use.

Prohibited Unacceptable	High risk Strict control	Limited risk Transparency	Minimal risk Free use

Level	Criteria	Example measures

Prohibited	Practices prohibited by Article 5 of the EU AI Act, and uses that directly lead to illegal acts	Not used for any reason (see Chapter 6)
High risk	Uses listed in Annex III of the EU AI Act, such as hiring, worker evaluation, credit scoring and education, and other decisions with legal or similarly significant effects on people	Approval by the AI Governance Committee, impact assessments, human oversight, log retention of at least six months
Limited risk	Chatbots facing customers, AI-generated content published outside the Company, deepfakes	Disclosure of AI use as required by Article 50, fact-checking, supervisor approval
Minimal risk	Internal drafts, summaries, search, translation, code completion	Only the data rules in Chapter 7 apply; free to use

6. Usage Rules

6.1 Approved tools (allowlist)

Generative AI tools may be used for work only if they have been reviewed and approved by IT Department / Digital Transformation Office. Using unapproved tools for work is prohibited.

Editor's note: Replace this list with the tools your company actually contracts and approves. Product names change, so check the current names at each revision.

Approved tools (examples)
• ChatGPT Business / ChatGPT Enterprise (text generation) • Claude Team / Claude Enterprise, Claude Code (text and code generation) • Microsoft 365 Copilot (with access to internal data) / Microsoft 365 Copilot Chat (web-grounded) • Gemini (included in higher Google Workspace plans) / Gemini Enterprise • GitHub Copilot Business / GitHub Copilot Enterprise (code generation) • Adobe Firefly (Teams / Enterprise business plans, image generation) • Internal LLM platform (confidential data allowed, with conditions)

Category	Conditions of use
Text generation (business contract)	After confirming that the contract does not allow input data to be used for training, data classified B or lower may be entered (see Chapter 7)
Code generation	Generated code is adopted only after review, testing and a license check. For client assets, the contract terms take precedence
Image, audio and video generation	Use within the scope of a license that permits commercial use. Output that resembles real people or existing works is checked before publication
AI agent platforms	Chapter 10 applies in full. Operations subject to approval go through human approval
Internal LLM platform	Data up to classification A may be entered, provided that access rights and usage logs are managed

Reviewing a new tool usually takes about two weeks. The review is not skipped even when the matter is urgent.

Personal accounts must not be used for work; only enterprise plans contracted by the Company may be used.

Requests to add a tool

Anyone who wishes to introduce a new AI tool submits an "AI Use Request Form" (Appendix B). The tool is approved after a security review, a check of its data protection policy and a legal review of its terms of use.

6.2 Permitted uses

- Drafting internal documents, reports and emails
- Summarizing meeting minutes and extracting action items
- Supporting market research and information gathering
- Generating, reviewing and debugging program code
- Supporting the creation of internal FAQs and knowledge bases
- Supporting translation and the creation of multilingual content
- Supporting data analysis and visualization, and brainstorming
- Automating business processes (on approved agent platforms)

When AI is used for content published or sent outside the Company, the content must be fact-checked and approved by a supervisor.

6.3 Prohibited acts

The following acts are prohibited for any reason

1. Entering personal information or confidential information of customers or business partners into AI
2. Publishing or sending AI output outside the Company without fact-checking it
3. Using AI to imitate or reproduce works created by others
4. Creating deepfakes (fabricated audio, images or video)
5. Using AI to psychologically manipulate employees or exploit their emotions
6. Making hiring or personnel evaluation decisions based on AI output alone
7. Using unapproved AI tools for work
8. Intentionally bypassing safety features or guardrails (jailbreaking)
9. Generating malicious code, malware or attack tools
10. Granting an agent permissions such as writing to external systems or deleting data without approval

Practices prohibited by the EU AI Act

Article 5 of the EU AI Act prohibits the following practices regardless of the tool used. They are prohibited in the Company as well, even where the list above does not mention them.

- Inferring the emotions of employees at the workplace, or of learners in education, except for medical or safety reasons
- Evaluating or classifying people by their social behavior or personal traits in a way that leads to unjustified detrimental treatment (social scoring)
- Using subliminal, manipulative or deceptive techniques, or exploiting a person's age, disability or social or economic situation, in ways that cause significant harm
- Categorizing people by biometric data to infer their race, political opinions, trade union membership, religious beliefs, sex life or sexual orientation
- Building facial recognition databases by untargeted scraping of facial images from the internet or CCTV footage
- Generating or manipulating sexual or intimate images of real people without their consent (prohibited from December 2, 2026)

6.4 Dealing with the use of unapproved tools (shadow AI)

A prohibition alone does not stop the use of unapproved tools. The Company provides both a mechanism for detecting such use and a contact point where it is easy to report it.

Means	Details
Network log review	Make access to generative AI services visible through a proxy, CASB or similar tool, and detect the use of unapproved services
Periodic inventory	Every six months, ask each department to report the AI tools it uses, and update the inventory sheet in Appendix E
Expense and purchasing checks	Check whether individual or departmental charges for AI services appear in expense claims
Self-reporting	Anyone who has already used an unapproved tool is not disciplined if they report it within the reporting period. Make the difference from cases that are concealed until an incident occurs clear

Editor's note: Keep only the means your company can actually carry out. Writing down means you will not implement creates rules that nobody follows. Whether to grant immunity for self-reporting is a management decision.

6.5 How to use AI by task, and what not to do

Editor's note: The more concretely you describe the situations where people hesitate, the better the rules are followed. Replace these rows with tasks that are common in your company.

Task	How to proceed	What not to do
Summarizing meeting minutes	Mask names and similar details in the recording or transcript before entering it, and have the person in charge check the facts after summarizing	Pasting text that still contains customer names or undisclosed deal terms into a personal plan
Creating materials and proposals	Have AI produce an outline and a first draft, and verify figures and proper nouns against the original sources	Putting market sizes or other companies' case studies cited by AI into a proposal without checking them
Translation	Translate public information and internal documents with approved tools, and have contracts and similar documents checked by an expert	Sending a contract or legal document to the other party as translated by AI
Generating and reviewing code	Use approved code generation tools, and adopt generated code only after review, testing and a license check	Entering code that contains client assets or credentials into external tools
Emails and inquiries	Have AI write a draft, and have the person in charge check the content and recipients before sending	Sending an AI draft to a customer without reading it
Research and information gathering	Have AI output source URLs as well, and check them against primary sources	Making decisions based on AI answers without checking the sources

6.6 Department-specific rules, personal devices and external communications

- **Additional department rules** Departments that handle highly sensitive information, such as HR, accounting, legal and public relations, have their department heads set rules in addition to this Policy. The additional rules are submitted to AI Governance Committee as annexes to this Policy
- **Personal devices and personal AI accounts (BYOAI)** Personal devices and personally contracted AI services are not used for work.
- **Use in external communications** When AI is used for content communicated outside the Company, such as social media posts, blogs or presentation materials, the content is checked by a supervisor before publication. AI output is not used as is for content presented as the Company's own view
- **Hiring, HR and evaluation** Applicant information and materials for personnel evaluation are not entered into AI. Decisions on selection and evaluation are made by people

6.7 Additional company-specific rules

Editor's note: These are company-specific arrangements raised in the interviews.

- Drawings, structural calculations and ground survey data must be handled only on the internal AI platform
- When AI is used for property descriptions or the wording of contracts, a qualified person must check the content before it is given to customers
- AI-generated images of properties may be used only if they do not give an impression that differs from the actual property
- Construction plans and safety management decisions must not be made on AI output alone

7. Data Handling Rules

7.1 Classification levels

Editor's note: Adjust these levels to match the classification in your information asset management rules.

Level	Classification	Examples	AI input
S	Strictly Confidential	Business strategy, M&A, unpublished financial results, the entire customer database, credentials and encryption keys	✕ Not allowed
A	Confidential	Customer personal information, contracts, technical specifications, personnel evaluations, private source code	△ Conditional (internal LLM only)
B	Internal Use Only	Internal manuals, meeting minutes, business reports, ordinary internal emails	✓ Allowed (enterprise plan)
C	Public	Press releases, information published on the web, product catalogs	✓ Allowed

7.2 Input permission matrix

Editor's note: This matrix reflects the interview results. Fine-tune it to match your operations.

Data type	Internal LLM	External (enterprise plan)	External (personal plan)
Public information (C)	✓ Allowed	✓ Allowed	✓ Allowed
Internal documents (B)	✓ Allowed	✓ Allowed (after anonymization)	✕ Not allowed
Customer information (A)	△ Conditional	✕ Not allowed	✕ Not allowed
Strictly confidential information (S)	✕ Not allowed	✕ Not allowed	✕ Not allowed

Personal information (names, addresses, etc.)	△ Anonymized data only	✕ Not allowed	✕ Not allowed
Source code (private)	✓ Allowed	△ Approved tools only	✕ Not allowed

7.3 Handling personal data

Entering personal data into an AI tool is processing under the GDPR. Personal data may be entered into AI only if all of the following are met.

- The use has a legal basis under Article 6 of the GDPR and is compatible with the purpose for which the data was collected
- The data is limited to what is necessary for the task (data minimization). Names and identifiers are removed or pseudonymized wherever the task can be done without them
- The tool is approved for personal data, and the provider is bound by a data processing agreement that prohibits using the data to train its models (see Chapter 12)
- Special categories of personal data (such as health data) and data about criminal convictions are not entered, except on the internal LLM with the approval described in Chapter 12
- As a rule, personal data is not entered into anything other than the internal LLM

> **Caution with multimodal input**
>
> Images, audio and video can also contain personal data. Before using cameras or speech recognition, inform the people concerned and check the legal basis. Biometric data used to identify people is a special category under the GDPR. For physical AI, also consider deleting data that remains on devices.

7.4 Connecting AI to internal data (RAG and internal document search)

If a system that lets AI search and refer to internal documents (RAG) does not carry over the access rights of the original documents, information that users are not supposed to see can appear in answers. The following rules apply both when the system is built and while it is in operation.

Issue	Rules to follow
Inheriting access rights	Reflect each user's viewing rights in the search scope so that documents the user has no right to see do not appear in answers or sources. Do not mix restricted documents, such as HR or accounting documents, into a company-wide shared index

Selecting the data	Define in advance what the index will contain, and exclude documents classified S. Take an inventory of the documents to include according to the classification in Chapter 7
Freshness and sources	Show the source document and its last update date together with each answer. To prevent outdated rules from being cited as current, set up a process for reflecting updates and deletions in the index
Departing and transferred employees	Reflect changes in access rights on the search side as well, so that the index does not keep old permissions
Responsibility of the builders	A department that builds and provides RAG internally is not just a user but a provider of AI, and is responsible for its quality and safety under the management described in Chapter 18

Editor's note: If you have not introduced internal document search or RAG, delete this section. If you plan to introduce it, keep the section and name the department that will build it.

7.5 Retaining and accessing input and output logs

Item	Arrangement
What is retained	The user, date and time, tool used, summary of the input and where the output was used. For agents, also the operations executed and who approved them (see Chapter 10)
Retention period	90 days (1 year for high-risk work)
Who may access the logs	Limited to staff designated by IT Department / Digital Transformation Office and investigators during incident investigations
Purpose of access	Limited to incident investigations, audits and aggregating usage. Logs are not used for personnel evaluation or monitoring individuals
Notice to users	Users are informed, when they start using the tools, that logs are collected and for what purpose

Editor's note: Collecting logs amounts to monitoring employees, so this Policy states that logs will not be used for other purposes. In many Member States, introducing a tool that can monitor employees requires informing or consulting the works council or employee representatives, and in some, their agreement. Check this before log collection begins.

8. Quality Control of AI Output

8.1 Fact-checking process

Risk level	Verification requirements	Verified by
High risk	Every item is confirmed against primary sources. Review by Legal or the relevant specialist department is mandatory	Specialist department staff + supervisor
Limited risk	Figures, proper nouns and legal statements are checked against reliable sources	Person in charge + supervisor
Minimal risk	The output is read through to confirm there are no obvious factual errors	Person in charge

8.2 Approval flow for using AI output

8.3 Countermeasures against hallucination

- Always check figures, dates, personal names, names of laws and URLs in the output against the original source
- Use retrieval-augmented generation (RAG) so that answers are grounded in reliable internal data sources
- Design prompts that ask the AI to state how confident it is
- Cross-checking the output with more than one model is also effective
- If an error leads to business damage, report it as an incident (see Chapter 20)

9. Copyright and Intellectual Property

9.1 Copyright in AI output

Under EU copyright law, a work is protected only if it is the author's own intellectual creation, reflecting free and creative choices made by a human. Output generated by AI without such human choices is generally not protected, and the Company cannot stop others from using it. Where a person makes creative choices, for example by substantially selecting, arranging or editing the output, the result may be protected to the extent of that contribution. Who owns works created by employees depends on national law and the employment contract.

9.2 Respecting the copyrights of others

- Do not use prompts intended to reproduce or imitate a specific protected work, artist or brand
- Before publishing, check that the output is not similar to existing copyrighted works
- For images, audio and video used commercially, check the terms of the commercial license
- Do not enter third-party works (articles, reports, images) into AI tools beyond what the license allows
- Consult Legal in advance before registering AI-generated material as a trademark or design

9.3 Text and data mining, and the Company's own content

EU law allows text and data mining of lawfully accessible content, including for training AI, unless the rights holder has reserved this right in an appropriate manner, such as machine-readable means for content published online (Article 4 of Directive (EU) 2019/790). Providers of general-purpose AI models must have a policy to respect such reservations (Article 53 of the EU AI Act).

- Data entered into external AI services may be used to train their models. Use only services whose contract excludes training on the Company's input
- Decide, with Legal and Marketing, whether **The Company** reserves its rights over content on its own websites against text and data mining, and how the reservation is expressed

9.4 Trade secrets

Know-how and business information are protected as trade secrets only while the Company takes reasonable steps to keep them secret (Directive (EU) 2016/943). Entering them into AI tools that are not approved for confidential data weakens the Company's position that it took such steps. The data rules in Chapter 7 are part of those steps.

9.5 Intellectual property in AI-generated code

- Generated code may carry open-source license conditions. Check license compatibility
- Do not adopt generated code as is. Put it through review and testing first
- If code being considered for a patent application contains AI-generated parts, consult the Intellectual Property Department

9.6 Labeling AI-generated content

- Images, audio or video that resemble real people, places or events and could falsely appear authentic (deepfakes) must be disclosed as artificially generated or manipulated (Article 50(4) of the EU AI Act). For evidently artistic, satirical or fictional works, the disclosure may be limited so that it does not spoil the work
- Text generated or manipulated by AI and published to inform the public on matters of public interest must be disclosed, unless a person has reviewed it and the Company takes editorial responsibility for it
- Do not remove watermarks or other machine-readable marks that tools add to AI-generated content

10. Rules for AI Agents

> **Basis for this chapter**
> The EU AI Act does not treat AI agents as a separate category. An agent is an AI system, and it becomes high-risk when it makes decisions in an area listed in Annex III, such as screening job applicants. Human oversight, least privilege and logging of operations are required for high-risk uses by Articles 14 and 26, and for any agent that touches personal data by the GDPR's duty to keep processing secure (Article 32).

10.1 Classification of AI agents

Type	Autonomy	Examples	Approval level
Task execution	Low	Drafting emails, routine reports	Department head
Workflow	Medium	Automating collection, analysis and report generation	Department head + IT Department
Autonomous decision-making	High	Automated customer support, purchasing decisions, operations automation	AI Governance Committee
Multi-agent	Highest	Agentic AI, automation of complex business processes	AI Governance Committee + senior management

10.2 Human-in-the-Loop (operations that require prior approval)

Critical operation	Approval method
Sending emails or messages outside the Company	Review of the draft before sending + approval
Deleting files or data	Display of the list of targets + explicit approval
Operations involving money (orders, transfers, purchases)	Display of the amount and counterparty + two-step approval
Applying changes to production environments	Display of the diff + approval + automatic rollback
Sending data to external APIs or services	Prior check of the destination and data + approval
Changing access permissions	Display of the changes + administrator approval

10.3 Permission design and emergency stop

- Grant only the minimum permissions needed to carry out the task

- Restrict connected tools and systems with an allowlist
- Separate read-only permissions from write permissions
- Give every agent an immediate stop function, and prepare procedures for restoring state after a stop
- Put in place a mechanism that stops an agent automatically when it detects unexpected repeated execution, mass access or communication with unexpected destinations
- Give the authority to stop agents to operations staff and above, and make it known in advance who can stop them
- Review granted permissions every quarter

10.4 Traceability requirements

Item	Details
When	Timestamp of the operation (ISO 8601 / UTC)
Who	Executing user ID + agent ID
What	Target of the operation and input (structured log)
Result	Success, failure or error details
Basis	Policy or rule applied (policy ID reference)

Logs must be retained for 90 days (1 year for high-risk work).

11. Security

11.1 Access management

- Manage accounts for AI services with single sign-on (SSO) or multi-factor authentication
- Disable the accounts of employees who leave or change roles immediately
- Manage API keys in a secrets manager and never write them directly in code
- Centralize usage logs and monitor them for suspicious access

11.2 Countermeasures against prompt injection

- Keep user input clearly separate from system prompts
- When external documents are retrieved through RAG, apply filters that detect malicious instructions embedded in them
- Limit the operations an agent accepts with an allowlist
- Sanitize output to prevent it from being executed as code or passed directly into SQL
- Run red-team exercises regularly

11.3 Security standards for external connections

Connected to	Security requirements
External APIs	TLS 1.2 or higher, API keys rotated at least every 90 days, encryption of transmitted data
Cloud storage	Access control lists, encryption at rest, audit logging enabled
Internal systems	Access via VPN or private network, API connections with least privilege
Open-source models	Confirm the developer is trustworthy, verify the download source, detect tampering

> **Security risks specific to AI agents**
>
> Agents connect to external systems on their own and carry out operations. Because an agent that receives malicious instructions can become a channel for sending internal data outside the Company, connections must be kept to the minimum necessary, and any new connection requires prior approval. Put in place a mechanism that detects unexpected communication or mass access, and stop the agent automatically when an anomaly occurs (see Chapter 10).

11.4 Checking models and providers

- At adoption, review technical documents such as model cards and system cards to understand the intended uses, known limitations and evaluation results

- Check the provider, the provenance of the training data, and the transparency of subcontracting and the supply chain
- When running a publicly released model in-house, verify that the download source is legitimate and that the model has not been tampered with
- For services that update or switch models automatically, design operations and testing on the assumption that the model's behavior may change

12. Privacy Protection

12.1 Principles

- Comply with the GDPR and the data protection law of each Member State in which the Company operates
- Process personal data in AI tools only for specified purposes and with a legal basis under Article 6 (for example, performance of a contract or legitimate interests). Where legitimate interests are the basis, record the balancing test
- Keep input to AI to the minimum necessary (data minimization), and pseudonymize or anonymize information that identifies individuals before entering it. Data counts as anonymous only if no one can single out a person, link records to a person or infer information about a person from it; pseudonymized data remains personal data
- Do not enter special categories of personal data (Article 9) into AI tools other than the internal LLM, and only with an exception under Article 9(2) confirmed by the data protection officer
- Tell people in the privacy notice when their data is processed with AI (Articles 13 and 14)
- If the output contains personal data, handle it with the same care as the input. Output about a person that is wrong is still personal data, and the person may ask for it to be corrected

12.2 Data protection impact assessment (DPIA)

A DPIA under Article 35 of the GDPR must be carried out before use begins where processing is likely to result in a high risk to people. In this Policy, a DPIA is required if any of the following applies, and the results are submitted to AI Governance Committee and the data protection officer.

- AI is used to evaluate, select or score customers, applicants or employees
- Data containing personal data will be processed by AI on an ongoing basis or on a large scale
- Data such as camera images, voice or location will be collected in ways the individuals may not notice
- Special categories of personal data or data about vulnerable people, such as children or patients, will be processed
- A high-risk AI system will be used. The provider's instructions for use feed into the DPIA (Article 26(9) of the EU AI Act)

Step	Details
1. Describe the processing	List the types of personal data, where it comes from, the purpose and legal basis, the retention period, the people involved and the AI tools and providers used
2. Assess necessity	Consider whether the personal data is really needed for the purpose, and whether a less intrusive method exists
3. Assess the risks	Assess the harm that could arise for individuals (incorrect decisions, discrimination, leaks, unexpected uses), and how severe and how likely it is

4. Decide on measures	Decide on measures that match the remaining risk, such as pseudonymization, access restrictions, shorter retention or human review. If high risk remains, consult the supervisory authority before starting (Article 36)
5. Record and review	Record the results of the assessment, and review them when the processing changes and at least once a year

12.3 Automated decisions

Decisions that have legal or similarly significant effects on people, such as rejecting a job applicant or a credit application, are not made solely by AI. A person with the authority and competence to change the outcome reviews each such decision (Article 22 of the GDPR). A score produced by AI can itself count as such a decision when it largely determines what someone else decides, as the Court of Justice held for credit scores in 2023. People may also ask how the decision was reached, and the Company must explain the procedure and principles actually applied to them. People affected by decisions based on high-risk AI may ask for an explanation of the role the AI played (Article 86 of the EU AI Act).

12.4 AI providers as processors

- Use only AI services covered by a data processing agreement under Article 28 of the GDPR that prohibits the provider from using the Company's data to train its models
- Check the list of sub-processors and where data is stored and processed
- Transfers of personal data outside the EEA need a basis under Chapter V of the GDPR, such as an adequacy decision or standard contractual clauses (see Chapter 15)

12.5 Requests from individuals

Requests for access, correction or erasure of personal data processed with AI are handled by the data protection officer within one month (Article 12(3)). Because data entered into an external AI service may be hard to locate or delete afterwards, the data rules in Chapter 7 keep personal data out of such services in the first place.

Editor's note: If **The Company** has not appointed a data protection officer, replace "data protection officer" with the person responsible for data protection. The risk assessment sheet in Appendix D can be reused to record a DPIA.

13. Risk Management

13.1 Risk classification

Category	Subcategory	Examples of risk
Technical risks	Training and input stage	Data poisoning attacks, unauthorized data mixed into training data
	Output stage	Biased or discriminatory output, inconsistent output, hallucination
	Post-deployment stage	Black-box behavior, insufficient explanation of decisions
Social risks	Ethics and law	Improper handling of personal information, accidents, discrimination, over-reliance, misuse
	Economic activity	Infringement of intellectual property, financial loss, leaks of confidential information, impact on employment
	Information space	Spread of false or misleading information, filter bubbles, loss of diversity
	Environment	Energy consumption and environmental impact

13.2 Risk assessment procedure

> **Governance that enables AI use**
>
> Excessive risk measures raise costs and discourage the use of AI. Allocate limited resources effectively in light of the nature and likelihood of each risk.

13.3 Risk mitigation measures

Risk area	Mitigation measures
Hallucination	RAG, fact-checking (see Chapter 8), adjusting the temperature parameter

Leaks of confidential information	Strict data classification (see Chapter 7), data loss prevention (DLP) tools, input filters
Bias and discrimination	Diversity testing of output, cross-checking, regular bias audits
Prompt injection	Sanitizing input and output, red-team exercises (see Chapter 11)
Over-reliance	Training (see Chapter 19), a culture in which a person always makes the final check
Copyright infringement	Checking commercial licenses, similarity checks (see Chapter 9)
Runaway agents or erroneous operations	Least privilege, limiting the operations that require approval, automatic stop on anomaly detection, logging of operations (see Chapter 10)
Defects in generated code	Mandatory review, static analysis, adequate test coverage, checks on dependencies and their licenses
Use of unapproved tools	Inspecting network logs, regular inventories, a channel for reporting (see Chapter 6)

13.4 Classifying a use under the EU AI Act

Before a new AI use starts, the applicant classifies it with the following questions, and the AI Governance Committee confirms the classification.

Question	If yes
Is it one of the practices prohibited by Article 5?	Not allowed (see Chapter 6)
Will it be used in recruitment or selection, decisions on promotion, dismissal or task allocation, or monitoring and evaluating workers?	High risk (Annex III, point 4)
Will it assess the creditworthiness of individuals, or price life or health insurance?	High risk (Annex III, point 5). A fundamental rights impact assessment is also required (Article 27)
Will it decide admission, evaluate learning outcomes or monitor learners in education?	High risk (Annex III, point 3)

Is it a safety component in the management of critical infrastructure, or of a product such as a machine or medical device?	High risk (Annex III, point 2, or Annex I)
Does it interact with people, or generate content published outside the Company?	Limited risk (Article 50 transparency)
None of the above	Minimal risk

13.5 Duties when the Company uses high-risk AI

The deployer duties for high-risk AI listed in Annex III apply from December 2, 2027. Uses that will continue beyond that date are prepared for now.

- Use the system according to the provider's instructions for use, and assign human oversight to people with the necessary competence, training and authority (Article 26)
- Make sure the input data under the Company's control is relevant and sufficiently representative for the purpose
- Monitor operation. If the system presents a risk, suspend its use and inform the provider and the market surveillance authority
- Keep the logs the system generates for at least six months
- Before using high-risk AI at the workplace, inform workers' representatives and the workers affected
- Inform people when high-risk AI is used to make or support decisions about them

> **When the Company becomes a provider**
>
> The Company takes on the duties of a provider if it puts its name or trademark on a high-risk AI system, makes a substantial modification to one, or uses an AI system, including a general-purpose one, for a high-risk purpose it was not intended for (Article 25). Changing the purpose of an approved tool therefore needs a new application.

13.6 Application dates of the EU AI Act

Obligation	Applies from	Details
Prohibited practices and AI literacy (Articles 4 and 5)	February 2, 2025	Applies now. The prohibitions on non-consensual intimate imagery and child sexual abuse material added in 2026 apply from December 2, 2026
General-purpose AI models	August 2, 2025	Duties of model providers, such as technical documentation, a copyright policy and a public summary of training content

Transparency (Article 50) and enforcement by national authorities	August 2, 2026	Telling people they are interacting with AI, marking AI-generated content, disclosing deepfakes and certain AI-generated text. Generative AI systems placed on the market before August 2, 2026 must have machine-readable marking by December 2, 2026
High-risk AI listed in Annex III	December 2, 2027	Employment, education, credit scoring, life and health insurance pricing, critical infrastructure and other areas. Includes the deployer duties in Articles 26 and 27
High-risk AI in products (Annex I)	August 2, 2028	AI that is a safety component of products such as machinery and medical devices

Editor's note: These dates reflect the Digital Omnibus on AI (Regulation (EU) 2026/1744) as of the date this template was prepared. Check the current dates before distribution.

14. Industry-Specific Considerations

Editor's note: These are the points to note for your industry (currently: Construction and real estate). Delete anything that does not apply.

14.1 Points to note by industry

Construction and real estate

Area	Points to note
Drawings and cost estimation data	Drawings, estimation bases and quotes from subcontractors may qualify as trade secrets. They may be entered into external AI only through approved tools.
On-site safety management	AI hazard detection is only a supporting tool, and responsibility for safety checks remains with the site manager. Operate on the assumption that detection will sometimes miss hazards.
Public works and tenders	Handle tender information as required by the contracting authority. Do not enter information about cost estimates or bid prices into external AI.

14.2 Additional controls common to all industries

- In regulated industries, sector-specific EU and national rules, supervisory guidance and industry self-regulation take precedence over this Policy
- Do not repurpose data entrusted by customers for the Company's own use of AI
- High-risk uses must go through a risk assessment (see Chapter 13) and any required approval before adoption

15. Regulations Outside the EU and Global Operations

Editor's note: If The Company has no offices, customers or data processing outside the EU and the EEA, you may delete this entire chapter.

15.1 Where the EU AI Act reaches

The EU AI Act applies to deployers established in the EU, and also to providers and deployers outside the EU where the output of their AI system is used in the EU. Group companies outside the EU that build or use AI affecting people in the EU are therefore in scope as well.

15.2 Developments in other countries and regions

Region	Overview
United Kingdom	No comprehensive AI law. The UK GDPR and the Data (Use and Access) Act 2025 govern personal data, and sector regulators apply AI principles within their own areas. EU AI Act duties may still apply when output is used in the EU.
United States	No comprehensive federal regulation. Voluntary frameworks such as the NIST AI RMF and state laws (California, Colorado and others) lead the way, and obligations differ from state to state.
China	Rules on generative AI services require algorithm filing, labeling of AI output and content controls. When providing services in China, check the local requirements case by case.
Japan and Korea	Japan's AI Promotion Act (2025) sets principles without penalties. Korea's AI Basic Act applies from January 2026 and includes duties for high-impact and generative AI.

15.3 Transfers of personal data outside the EEA

- Using an AI service that processes data outside the EEA is a transfer under Chapter V of the GDPR. It needs an adequacy decision, standard contractual clauses (SCCs) or another transfer tool
- For US providers, check whether the provider is certified under the EU-US Data Privacy Framework. If it is not, use SCCs and a transfer impact assessment. Because the Framework has been challenged before the Court of Justice, keep SCCs in the contract as a fallback
- Specify in the contract the countries where data is stored and processed, and where sub-processors are located. Prefer services that let you keep data in an EU region

15.4 Application to offices and group companies outside the EU

- This Policy is the minimum standard for the whole group. Offices in each country add stricter rules where local law requires them

- If local law conflicts with this Policy, local law takes precedence, and the conflict must be reported to the managing department named in Chapter 18
- Additional rules for each office are managed as annexes to this Policy

16. Procurement, Contracts and Vendor Management

Editor's note: These are the review criteria for adopting AI tools and services. Adjust them to fit how your IT, Procurement and Legal departments work.

16.1 Checklist for reviewing new tools

Area	What to check
Data handling	Whether input data is used for training, whether an opt-out is available, where and for how long data is stored, how data is deleted
Security	Third-party certification (ISO 27001, SOC 2 or similar), encryption, access control, vulnerability management
Compliance	Applicable laws, handling of personal information and cross-border transfers, legal review of the terms of service
Business continuity	SLA, support arrangements, vendor lock-in, data migration on exit
Model authenticity	Provenance of the provider and model, transparency of subcontracting and the supply chain

16.2 Points to specify in contracts

- A data processing agreement under Article 28 of the GDPR, covering the subject matter, duration, sub-processors, security measures, assistance with requests from individuals and deletion at the end of the service
- That input and output data are not used to train the provider's models, and who owns them
- The countries where data is stored and processed, and the transfer tool used for any transfer outside the EEA
- For high-risk AI, the instructions for use and the information the Company needs as a deployer, including the provider's cooperation when the Company changes or rebrands the system (Article 25(2) of the EU AI Act)
- Notification of security incidents and personal data breaches without undue delay, so that the Company can meet the 72-hour deadline in Chapter 20
- Return or deletion of data when the service ends or the contract is terminated

Editor's note: Financial entities also enter AI services in the register of ICT third-party arrangements under DORA, and include the contract terms that DORA requires.

17. Developing Systems with Built-in AI

Editor's note: If no department in your company builds systems or features that use AI, delete this entire chapter. If you outsource development, the chapter is still worth keeping, on the basis that you require the contractor to meet the same standards.

The chapters so far set rules for the Company as a user of generative AI. This chapter sets the additional rules that apply when the Company builds AI into its own products or internal systems and provides them to others. As a provider, the Company takes on responsibilities toward the people who use those systems.

17.1 Planning and design

- Decide at the design stage which tasks the AI performs and which decisions people make. Write the points where human judgment is required into the specification
- Identify the harm users could suffer from incorrect output. Uses with significant impact must go through a risk assessment (Chapter 13)
- If personal data is involved, conduct an impact assessment (Chapter 12) in parallel with the design work
- For uses that could disadvantage people, such as hiring, credit decisions or performance evaluations, build explainability and a means of appeal into the design

17.2 Implementation

- Keep the system prompt separate from user input, and make sure the system does not follow instructions embedded in documents imported from outside (Chapter 11)
- Grant the AI only the minimum permissions it needs. Writing to, deleting from or sending data to external systems must require approval or an explicit user action
- Log inputs, outputs, the model used and its version so that behavior can be traced afterward
- Manage API keys and credentials with a secrets management system. Never write them directly into prompts or code

17.3 Evaluation and release

Aspect	What to check
Accuracy	Evaluate output quality on realistic business data, and decide in advance what error rate is acceptable
Safety	Test for inappropriate output, leakage of confidential information and hijacking of instructions
Fairness	Check representative cases for output bias based on personal attributes
Regression	When the model or prompts change, rerun the same tests to confirm that quality has not dropped

Disclosure to users	Tell users, in a way they can understand, that AI is used, what its limitations are and whom to contact

17.4 Operations and change management

- Model updates by the provider can change the system's behavior. Define a procedure for detecting such changes and re-evaluating the system
- Set up a channel for users to report errors, and connect it to the reporting route in Chapter 20
- Define the criteria and procedure for suspending the service, so that it can be stopped promptly if a serious error comes to light
- Adding a connection to an external tool or data source requires prior approval. Maintain a list of all connections

17.5 Duties as a provider under the EU AI Act

- If a system the Company builds is high-risk, the requirements for providers apply from December 2, 2027 for Annex III systems: risk management, data governance, technical documentation, logging, transparency to deployers, human oversight, accuracy and robustness, a quality management system and conformity assessment
- A chatbot or other system that interacts with people must be designed so that people know they are dealing with AI, and generated audio, images, video and text must be marked in a machine-readable way (Article 50(1) and (2))
- If the system processes personal data, build in data protection by design and by default (Article 25 of the GDPR)
- Products with digital elements sold in the EU are also subject to the Cyber Resilience Act, whose reporting duties apply from September 11, 2026 and main obligations from December 11, 2027. Coordinate with the product security team
- Under the new Product Liability Directive (EU) 2024/2853, software and AI systems placed on the market from December 9, 2026 are products, and the Company can be liable for damage caused by their defects without proof of fault

18. Organization and Governance

18.1 Governance structure

Editor's note: Replace the department names and job titles to match your actual organization.

Role	Assigned to	Responsibilities
Chief AI Officer (CAIO)	Executive Officer in charge of IT	Sets AI strategy, approves this policy and holds final authority over risk decisions
AI Governance Committee	Representatives of IT, Legal, HR, Digital Transformation and each business unit	Operates and revises this policy, reviews risk assessments and sets the incident response policy
AI Champions	At least one per department	Promote AI use within the department, raise awareness of compliance and report on usage
IT Department / Digital Transformation Office	IT	Reviews and deploys tools, manages security, monitors logs and provides technical support
Legal Department	Legal	Assesses legal risk, reviews terms of service and handles legal matters during incidents
Data Protection Officer	DPO (where appointed)	Advises on data protection, reviews DPIAs and handles requests from individuals and contacts with the supervisory authority
Internal Audit Department	Audit	Audits compliance regularly and recommends improvements

18.2 Responsibilities of senior management

- Establish the AI governance policy and communicate it company-wide
- Oversee the activities of the AI Governance Committee and receive a report from it every quarter
- Make prompt decisions and give response instructions when an incident occurs
- Make management decisions that balance AI adoption and risk management
- Decide on the investments needed for AI use, such as tools, platforms and training

18.3 Operation of the AI Governance Committee

- Meeting frequency: Monthly regular meeting + ad hoc meetings as needed
- Agenda: review of new uses, review of assessment results, incident reports and proposed revisions
- Prepare minutes and report them to senior management

18.4 From application to approval

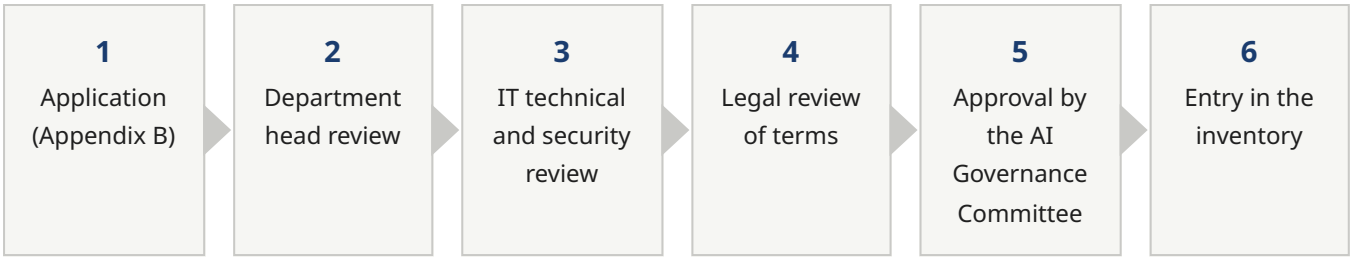

Item	Arrangement
Standard review period	Usually about two weeks from receipt of the application
High-risk uses	Attach the results of the risk assessment (Chapter 13) and the impact assessment. The AI Governance Committee reviews each case individually
If the application is rejected	The applicant is told the reasons in writing and may reapply once the conditions are met
After approval	Enter the tool in the inventory in Appendix E, and recheck it when the contract is renewed

18.5 Monitoring usage

The AI Governance Committee compiles the following metrics every quarter and reports them to senior management. The figures are used to check both whether AI adoption is progressing and whether the agreed rules are being followed.

Metric	What it shows
License utilization rate	The share of purchased licenses actually in use. A low rate points to insufficient training or too few identified use cases
Users and frequency of use by department	Whether use is concentrated in a few departments
New tool applications and approval rate	Demand from the front line and bottlenecks in the review process
Incidents and near misses	How effective the controls are. Very few reports may mean that reporting is difficult
Training completion rate and test results	Progress of the training in Chapter 19

Unapproved tools detected	The spread of shadow AI (Chapter 6)

Editor's note: You do not need to measure everything from the start. It is easier to keep going if you begin with the metrics you can already collect and decide first where they will be reported.

18.6 Audits

- The Internal Audit Department audits compliance with this policy at least once a year
- Audits cover the use of tools other than approved tools, the actual entry of confidential information, log retention, records of applications and approvals, and training completion
- To prepare for audits, keep application forms and approval records, risk assessment results, the inventory, incident reports and training records as evidence. Templates for all of them are in the appendix
- For each audit finding, prepare a corrective action plan. The AI Governance Committee confirms that the corrective action is complete
- If the Company seeks external certification, such as ISO/IEC 42001, map this policy to the certification requirements

18.7 Transferred employees, departing employees and contractors

Subject	Arrangement
Transferred employees	Review access rights at the time of transfer. Do not leave the employee able to search or view confidential documents of their former department (Chapter 7)
Departing employees	Disable the account on the last working day. Confirm how business data entered into or stored in AI tools is handled, and obtain a written pledge that the employee will not take it with them
Contractors	Make compliance with this policy a condition of the contract. In particular, state the prohibited acts in Chapter 6, the data handling rules in Chapter 7 and the reporting obligations in Chapter 20
Subcontracting	Requires prior approval. Subcontractors are subject to the same obligations
Audit rights	The contract must allow the Company to request reports on the contractor's compliance and to verify it when necessary
End of contract	Suspend accounts provided by the Company, have entrusted data returned or deleted, and obtain a certificate of deletion

19. Education and AI Literacy

Audience	Content	Frequency	Format
All employees	AI literacy basics, risk awareness and prohibited acts	On joining and once a year	E-learning and a comprehension test
Each department	Job-specific AI skills and prompt design	Every six months	Workshop
Managers	Governance, risk judgment and initial incident response	Once a year	Classroom training
Engineers	Secure development, agent design and building RAG systems	Every quarter	Hands-on training
AI Governance Committee members	Regulatory trends, overseas regulations and leading practices	Every quarter	Study sessions

> **Preventing overreliance**
>
> Training also covers the risk of psychological dependence on AI and excessive trust in it. Maintain a culture in which people make the final decision, and build the critical thinking needed to avoid accepting AI output uncritically.

> **AI literacy under the EU AI Act**
>
> Article 4 of the EU AI Act requires providers and deployers to take measures to support the AI literacy of their staff and of others who operate or use AI systems on their behalf. It does not require any particular level for each person, or a certificate. The Company meets this duty with the training above, matched to each person's role and the AI they use, and keeps the records in Appendix F.

20. Incident Response

20.1 Definition of an incident

- Confidential or personal information was entered into an AI service (data leak)
- Damage resulted from a decision or action based on AI output
- AI output was alleged to infringe copyright
- An agent performed an unintended operation, such as placing a wrong order, sending to the wrong recipient or deleting data
- An attack on an AI system was detected
- AI output contained discriminatory or harmful content

20.2 Response flow

Step	Owner	Action
Discovery	The person who discovers it	Stop the use immediately and trigger the emergency stop for agents. Notify your manager and IT
Initial response	IT and the person who discovered it	Determine the scope of damage, preserve evidence and prevent the damage from spreading
Report	Manager → AI Governance Committee	Prepare the incident report (Appendix C). Report serious incidents to senior management immediately
Root cause analysis	AI Governance Committee	Identify the root cause, confirm the scope of impact and decide whether notifications are required
Prevention	AI Governance Committee and related departments	Improve operations, add training, and report externally or disclose publicly where necessary

20.3 Notification deadlines under EU law

Event	Deadline	What to do

Personal data breach	Within 72 hours of becoming aware	Notify the supervisory authority unless the breach is unlikely to result in a risk to people (Article 33 of the GDPR). If the risk is high, also inform the people affected without undue delay (Article 34). Record every breach, including those not notified
Serious incident involving high-risk AI	Immediately	As deployer, inform the provider first, then the importer or distributor and the market surveillance authority (Article 26(5) of the EU AI Act)
Significant incident at an entity in scope of NIS2	Early warning within 24 hours, notification within 72 hours	Report to the national CSIRT or competent authority, followed by a final report within one month
Major ICT-related incident at a financial entity	As set by DORA	Report to the competent authority using the DORA templates

Editor's note: Keep only the rows that apply to **The Company**. The person who decides whether to notify the supervisory authority is usually the data protection officer.

21. Consequences of Violations

Severity	Examples	Measures
Minor	Temporary use of an unapproved tool with no confidential input; skipping a procedure with no actual harm	Verbal warning and retraining
Moderate	Entering internal-only data into an unapproved tool; breaching a prohibited act with minor actual harm	Written warning, temporary suspension of access, and mandatory retraining with a passing test result
Serious	Entering confidential or personal information into an external AI service, deliberate jailbreaking, creating deepfakes, or causing serious damage	Disciplinary action under the Company's disciplinary rules and consideration of legal action

The AI Governance Committee determines whether a violation has occurred. Disciplinary action is decided under the Company's disciplinary rules and applicable employment law after consultation with HR. Repeated violations are subject to the next level of measures, even if they are unintentional.

22. Cost and Environmental Considerations

Editor's note: Align the environmental section with your sustainability reporting. You may also keep only the cost management section.

22.1 Cost management

Item	Arrangement
Visibility of costs	Aggregate license fees and pay-as-you-go API charges by department, and include them in the quarterly report described in Chapter 18
Pay-as-you-go charges	Set a spending cap or request cap for API use, with a notification when it is exceeded. Continuous execution by agents is especially likely to exceed budgets
Choosing models	Do not routinely use models more powerful than the task requires. Lightweight models are often enough for tasks such as summarizing and classification
Idle licenses	Identify licenses that have not been used for three months or more during the inventory, and reassign or cancel them
Cost-effectiveness	Track the time and work saved by each department, and use the results in the next year's investment decisions

22.2 Environmental considerations

- Large-scale training and inference consume a lot of electricity. Avoid unnecessary bulk runs, and use sample data for trials
- For tasks that repeat the same processing, reduce the number of runs by reusing results or caching them
- If **the Company**'s sustainability policy requires reporting of environmental impact, agree with the responsible department on how AI use is accounted for

23. Periodic Review and Revision Cycle

Type	Frequency	Trigger
Scheduled revision	Every six months	Periodic review on a set schedule
Ad hoc revision	As needed	Major legal changes, newly emerging risks, technological shifts or major changes in the business environment

The AI Governance Committee drafts each revision, which is announced company-wide after approval by senior management. The revision history is kept under version control so that previous versions remain available for reference.

24. Laws and Guidelines Referenced

Editor's note: Laws and guidance continue to be revised. The dates in this table are as of the time this document was prepared. Check the latest versions before distribution, and add the national laws of the Member States in which you operate.

Name	Issued by	Notes
Artificial Intelligence Act (Regulation (EU) 2024/1689)	European Parliament and Council	In force since August 1, 2024. Applied in stages until August 2, 2028
Digital Omnibus on AI (Regulation (EU) 2026/1744)	European Parliament and Council	In force since July 27, 2026. Postponed high-risk obligations and revised Article 4
Guidelines on prohibited AI practices	European Commission	February 4, 2025
Guidelines on the definition of an AI system	European Commission	February 6, 2025
Guidelines on the transparency obligations under Article 50	European Commission	July 20, 2026
Code of Practice on transparency of AI-generated content	European Commission (AI Office)	June 10, 2026. Voluntary
AI literacy: questions and answers	European Commission	Updated October 20, 2026 to reflect the Digital Omnibus
General-Purpose AI Code of Practice	European Commission (AI Office)	July 10, 2025. For providers of general-purpose AI models
General Data Protection Regulation (Regulation (EU) 2016/679)	European Parliament and Council	Applies since May 25, 2018
Opinion 28/2024 on data protection aspects of AI models	European Data Protection Board (EDPB)	December 2024
Guidelines 02/2026 on anonymisation	European Data Protection Board (EDPB)	Version 1.0, July 2026. Public consultation until October 30, 2026

Directive on copyright in the Digital Single Market (Directive (EU) 2019/790)	European Parliament and Council	Articles 3 and 4: text and data mining
Trade Secrets Directive (Directive (EU) 2016/943)	European Parliament and Council	Reasonable steps to keep information secret
NIS2 Directive (Directive (EU) 2022/2555)	European Parliament and Council	Cybersecurity risk management and incident reporting for essential and important entities
Digital Operational Resilience Act (Regulation (EU) 2022/2554)	European Parliament and Council	Applies to financial entities since January 17, 2025
Product Liability Directive (Directive (EU) 2024/2853)	European Parliament and Council	Applies to products placed on the market from December 9, 2026, including software and AI systems
Cyber Resilience Act (Regulation (EU) 2024/2847)	European Parliament and Council	Reporting duties from September 11, 2026; main obligations from December 11, 2027
Ethics Guidelines for Trustworthy AI	High-Level Expert Group on AI	April 2019. Seven requirements used in Chapter 5
ISO/IEC 42001:2023	ISO/IEC	AI management systems

25. Frequently Asked Questions (FAQ)

Editor's note: Replace these with the questions your employees actually ask most often. People will refer to this section more.

Question	Answer
May I use a free personal ChatGPT account for work, just a little?	No. Use only the approved business editions. With personal editions, what you enter may be used for training, and there is no data processing agreement, so entering personal data would breach the GDPR.
I want AI to summarize the minutes of an internal meeting. May I enter them?	If the minutes are classified Internal Use Only (class B), you may enter them into an approved business edition. Remove or pseudonymize names and other personal data first.
May I use AI to screen job applications?	Only with the approval of the governance committee. Screening applicants is high-risk under the EU AI Act, and decisions about applicants must not be made by AI alone under the GDPR.
Do I have to label everything I write with AI's help?	No. Labeling is required for deepfakes, for chatbots that talk to people, and for AI-generated text published to inform the public that nobody has reviewed. Internal drafts you have checked and edited do not need a label.
May I use figures produced by AI in my materials as they are?	No. Always check numbers, dates, proper nouns and URLs against the original sources.
What should I do if I want to use a new AI tool?	Submit the application form in Appendix B. The tool can be used once it has been reviewed and approved.

26. Appendix: Fill-in Templates

Appendix A: Self-check list

A-1. Before using AI

- [ ] The tool is on the Company's list of approved tools (Chapter 6)
- [ ] I checked the classification of the input data (Chapter 7)
- [ ] I am not entering personal or confidential information, or it has been anonymized
- [ ] The purpose falls within the permitted scope of work
- [ ] I am signed in with an Enterprise account

A-2. When using AI output

- [ ] I fact-checked the output (Chapter 8)
- [ ] I checked numbers, dates, names of people, names of laws and URLs against the original sources
- [ ] I obtained approval appropriate to the risk before publishing externally
- [ ] Where disclosure of AI generation is required, I included it
- [ ] I confirmed that there is no concern about copyright infringement (Chapter 9)

A-3. When using AI agents

- [ ] I obtained approval appropriate to the agent's level of autonomy (Chapter 10)
- [ ] Human-in-the-Loop is implemented for critical operations
- [ ] Permissions follow the principle of least privilege
- [ ] An emergency stop is implemented
- [ ] The five traceability items are being recorded
- [ ] Connections to external systems are limited to the minimum necessary

Appendix B: AI use application form

Item	Entry
Date of application	
Applicant (name and department)	
Tool to be used	
Provider and contract type	
Purpose of use	

Classification of input data	
Department and number of users	
Risk level (self-assessment)	
Risk mitigation measures	
Approval (IT)	
Approval (Legal)	
Approval (AI Governance Committee)	

Appendix C: Incident report

Item	Entry
Date and time of report	
Reporter (name and department)	
Date and time of occurrence	
Incident type	Data leak / Incorrect output / Copyright / Agent malfunction / Attack detected / Other
Tool or system used	
Summary	
Scope of impact	
Initial response	
Cause (as far as known)	
Preventive measures (proposed)	
Response status	Initial response complete / Under investigation / Measures implemented / Closed

Appendix D: AI risk assessment sheet

Item	Entry
AI use or system assessed	
Purpose and scope of work	
Type and classification of input data	
Anticipated risks	
Impact (high / medium / low)	
Likelihood (high / medium / low)	
Overall risk level	Prohibited / High / Limited / Minimal
Mitigation measures	
Decision (approved / conditional / rejected)	
Assessor and date	

Appendix E: Inventory of approved AI tools

Tool	Purpose	Contract type	Permitted data	Date approved

Appendix F: Training records

Employee or department	Training	Date completed	Test result