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 Japanese law and Japanese government guidance. This policy was prepared with reference to the AI Guidelines for Business (Version 1.2) issued by the Ministry of Internal Affairs and Communications and the Ministry of Economy, Trade and Industry; the Act on the Promotion of Research, Development and Utilization of AI-Related Technologies (AI Promotion Act) and the Basic Plan for Artificial Intelligence of the Cabinet Office; the Digital Agency's Guideline on the Procurement and Utilization of Generative AI for the Evolution and Innovation of Public Administration (Version 2.0); the EU AI Act; and other sources.

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 reference laws in Chapter 1 and figures such as log retention periods and reporting deadlines are general benchmarks as of the time this document was prepared. 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. Overseas Regulations 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 laws, regulations and public guidelines.

Law or guideline	Issued by	Status
AI Guidelines for Business (Version 1.2)	Ministry of Internal Affairs and Communications (MIC) and Ministry of Economy, Trade and Industry (METI)	Published in March 2026. Covers physical AI and AI agents
Act on the Promotion of Research, Development and Utilization of AI-Related Technologies (AI Promotion Act)	Cabinet Office	Promulgated in June 2025 and fully in force since September 2025. Sets out principles and promotion measures (no penalties)

Basic Plan for Artificial Intelligence	Cabinet Office	Adopted by the Cabinet in December 2025. A statutory plan under the AI Promotion Act
Guideline on the Procurement and Utilization of Generative AI for the Evolution and Innovation of Public Administration (Version 2.0, DS-920)	Digital Agency	Issued in May 2025 and revised to Version 2.0 on June 12, 2026. Written for government agencies, but its procurement review criteria are also useful for private companies
Guidelines for Ensuring the Appropriateness of Research, Development and Utilization of AI-Related Technologies	Cabinet Office	Guidelines under the AI Promotion Act. Describe what developers and users should observe to ensure appropriate use
EU AI Act (Regulation (EU) 2024/1689)	EU	Entered into force in August 2024. Transparency obligations apply from August 2026. High-risk obligations were postponed to December 2027 or later by the Digital Omnibus (Regulation (EU) 2026/1744, in force since July 27, 2026). Applies to companies outside the EU
Act on the Protection of Personal Information (APPI)	Personal Information Protection Commission (PPC)	Amendment enacted on July 10, 2026 and promulgated on July 17, 2026 (Act No. 56 of 2026). Introduces a consent exception for AI development and statistical purposes and an administrative surcharge system. Penalty-related provisions take effect early, on January 17, 2027; the main provisions take effect on a date set by Cabinet Order, no later than July 16, 2028
Unfair Competition Prevention Act	Ministry of Economy, Trade and Industry (METI)	Protects trade secrets. Note that entering information into AI may cause it to lose its status as a managed secret

Copyright Act (Article 30-4, etc.)	Agency for Cultural Affairs	Published its views on AI training and generation (March 2024)
ISO/IEC 42001:2023	ISO/IEC	International standard for AI management systems

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

Key terms follow the definitions in the AI Guidelines for Business (Version 1.2).

Term	Definition
Generative AI	A general term for AI systems, built on LLMs and similar models, that generate text, images, audio, video, code and other content.
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.
Agentic AI	A system in which multiple AI agents cooperate autonomously to make decisions and take actions.
Physical AI	A system that captures the physical environment through sensors, processes it with AI, and turns the result into physical actions through actuators and similar devices.
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.
Fine-tuning	Additional training of a pre-trained model for a specific task.
Human-in-the-Loop	A mechanism that builds human involvement into AI judgments and actions.
Prompt injection	An attack that uses malicious input to bypass controls and cause unintended behavior. Includes indirect attacks.

5. Basic Principles

5.1 Ten common principles

The ten principles of the AI Guidelines for Business (Version 1.2) are interpreted and applied in the context of the Company's business.

No.	Principle	Interpretation and application
1	Human-centricity	AI output is reference information, and humans make the final decisions. AI is not used to manipulate decisions or exploit emotions.
2	Safety	Usage policies are set to prevent harm. The authenticity of models is verified.
3	Fairness	Bias is taken into account, and output is checked so that discriminatory results do not influence business decisions.
4	Privacy protection	Entering personal information and privacy-related data into AI is strictly restricted.
5	Security	Technical and organizational measures are taken against injection, data poisoning and similar threats. External integrations use the least privilege necessary.
6	Transparency	Content created with AI is labeled as such. Customers are informed in advance when AI is used to serve them.
7	Accountability	The organization is accountable for the results of AI use. Usage logs are preserved.
8	Education and literacy	All employees receive training. Excessive dependence is prevented and critical thinking is maintained.
9	Fair competition	The Company does not engage in unfair market dominance or data monopolization.
10	Innovation	Chilling effects from excessive regulation are avoided, and control is graded according to risk.

5.2 Risk-based approach

Rather than regulating all AI use uniformly, the Company prioritizes measures according to the severity and likelihood of risks.

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

Level	Criteria	Example measures

Prohibited	Directly leads to human rights violations or illegal acts	Social scoring, employee surveillance and emotional manipulation, impersonation
High risk	Decisions on HR, credit, legal matters, or medical care and safety	Final decision by a human, risk assessment, periodic audits, log preservation
Limited risk	Customer-facing content and materials published outside the Company	Disclosure of AI use, fact-checking, supervisor approval
Minimal risk	Internal drafts, summaries, search, translation, code completion	Only the ban on entering confidential data applies; 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

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.

- Residents' personal information must be handled only on the internal AI platform
- AI procurement must be reviewed in line with the Digital Agency's Guideline on the Procurement and Utilization of Generative AI for the Evolution and Innovation of Public Administration (Version 2.0)
- When AI is used to draft answers to assembly questions or official notices, the facts must still be checked against primary sources
- Responses to residents drafted with AI must be approved by the responsible section before they are sent

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 information

In accordance with the Act on the Protection of Personal Information (APPI), personal information may be entered into AI only if all of the following are met.

- The use is within the scope of the purpose of use
- The data is limited to the minimum necessary (data minimization)
- No unnecessary attribute information is attached
- The data has been anonymized or pseudonymized
- As a rule, input into anything other than the internal LLM is prohibited

> **Caution with multimodal input**
>
> Images, audio and video can also contain personal information. Systems that use cameras or speech recognition must notify the people concerned and obtain their consent. 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. Check within your company whether an explanation to the labor union or employee representative is required.

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 the current Copyright Act, content generated autonomously by AI is not protected by copyright. Where a human makes a creative contribution, however, the output may qualify as a copyrighted work depending on the extent of that contribution. As a rule, output produced in the course of business is treated as a work made for hire owned by the Company.

9.2 Respecting the copyrights of others

- Do not use prompts intended to imitate the "style" of a specific 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
- Consult Legal in advance before registering AI-generated material as a trademark or design

9.3 Company content as training data

Data entered into external AI services may be used to train their models. Enterprise contracts often guarantee an opt-out, but the terms of service must be checked.

9.4 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

10. Rules for AI Agents

> **Basis for this chapter**
> Version 1.2 of the AI Guidelines for Business added definitions, risks and points of attention for AI agents and physical AI. Human-in-the-Loop, least privilege and logging of operations have become de facto requirements.

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

- Comply with the Act on the Protection of Personal Information (APPI), the GDPR (when processing data in the EU) and all other applicable laws
- Keep input to AI to the minimum necessary (data minimization)
- Anonymize or pseudonymize information that can identify individuals before entering it
- If the output contains personal information, handle it appropriately and respect the privacy of the people concerned
- For systems that use cameras or speech recognition, notify people of the purpose of use and obtain their consent
- For physical AI, delete data left on devices regularly
- Conduct a privacy impact assessment (PIA) before starting any high-risk use of AI

12.1 Conducting impact assessments (PIA / DPIA)

If any of the following applies, an impact assessment must be conducted before use begins and the results submitted to AI Governance Committee.

- Data containing personal information will be processed by AI on an ongoing basis
- AI will be used to evaluate, select or score customers or employees
- Data such as camera images, voice or location will be collected in ways the individuals may not notice
- Special care-required personal information will be handled, or personal data in the EU will be processed (a DPIA under the GDPR may be required)

Step	Details
1. Identify the scope	List the types of personal data to be processed, where it comes from, the purpose of use, the retention period and the people involved
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 anonymization, access restrictions, shorter retention or human review
5. Record and review	Record the results of the assessment, and review them when the processing changes and at least once a year

Editor's note: The risk assessment sheet in Appendix D can be reused to record the assessment. If you use a separate form, write the name of your form in this section.

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)

14. Industry-Specific Considerations

Editor's note: These are the points to note for your industry (currently: Public sector and local government). Delete anything that does not apply.

14.1 Points to note by industry

Public sector and local government

Area	Points to note
Administrative and resident information	Comply with the Digital Agency's Guideline on the Procurement and Utilization of Generative AI for the Evolution and Innovation of Public Administration. Residents' personal information must not be entered into external AI.
Fairness and accountability	AI output must not be the basis for adverse administrative dispositions. Ensure that the decision process can be explained and that records are preserved.
Procurement and outsourcing	When procuring AI, specify in the contract the provenance of training data, subcontracting and cross-border data transfers.

14.2 Additional controls common to all industries

- In regulated industries, industry-specific laws, supervisory guidelines 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. Overseas Regulations and Global Operations

Editor's note: If you have no overseas offices, no overseas customers and no products for the EU, you may delete this entire chapter.

15.1 EU AI Act

The EU AI Act regulates AI according to four risk categories: unacceptable risk, high risk, limited risk and minimal risk. Because it can also apply to companies outside the EU when their AI affects users or markets in the EU, any company providing products or services to the EU must check which category applies to it. High-risk AI (used in recruitment, credit, education, critical infrastructure and similar areas) is subject to obligations covering risk management systems, data governance, human oversight, technical documentation and record-keeping, and conformity assessment.

Obligation	Applies from	Details
Prohibited AI practices	In effect (February 2025)	Bans on subliminal manipulation, social scoring, untargeted scraping of facial images and similar practices
Obligations for general-purpose AI (GPAI) models	In effect (August 2025)	Technical documentation, a public summary of training data, a copyright policy
Transparency obligations (Article 50)	August 2, 2026	Telling people they are interacting with AI, labeling AI-generated content, disclosing emotion recognition and biometric categorization
Machine-readable marking of AI output	December 2, 2026	Adding watermarks or similar marks to the output of generative AI
High risk (Annex III systems)	December 2, 2027	Stand-alone systems used in employment, education, access to essential services and similar areas. Postponed from the original date of August 2026
High risk (Annex I systems)	August 2, 2028	Systems embedded in products. Postponed from the original date of August 2027

> **Note that the high-risk obligations have been postponed**
>
> The Digital Omnibus adopted in July 2026 (Regulation (EU) 2026/1744, published in the Official Journal on July 24, 2026 and in force from July 27, 2026) pushed back the compliance deadlines for high-risk AI as shown above. The August 2, 2026 date for the transparency obligations, however, remains unchanged. The postponement gives more time to prepare and is not an exemption, so companies in scope must continue preparing for conformity assessment. Check the latest application dates before distributing this Policy.

15.2 Developments in major countries and regions

Region	Overview
United States	There is no comprehensive federal regulation, and a change in the administration's policy has shifted the direction toward promoting AI and easing regulation. In practice, voluntary frameworks such as the NIST AI RMF and state laws (California, Colorado and others) lead the way. Note that obligations differ from state to state.
EU	Comprehensive, cross-sector regulation through the AI Act, with extraterritorial reach. Applied in stages as described at the start of this chapter.
China	The Interim Measures for the Management of Generative AI Services and related rules require algorithm filing, labeling of AI output and content controls. When providing services in China, check the local requirements case by case.
United Kingdom	No comprehensive AI law. Existing regulators apply principles within their own sectors.
Others	Korea, Canada and other countries are developing AI legislation. Key issues are international interoperability and alignment with international norms such as the Hiroshima AI Process.

15.3 Cross-border data transfers

- Using an overseas AI service may count as a cross-border transfer of data. Check the level of protection in the destination country, and whether consent or a compliance framework is required under the APPI
- For data subject to the GDPR, secure a lawful basis for transfer, such as standard contractual clauses (SCCs)
- Specify in the contract the countries where data is stored and processed, and where subprocessors are located
- For data that must be kept in Japan (because of industry-specific laws or contractual restrictions), choose services that allow you to specify the region

15.4 Application to overseas offices and group companies

- 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

- Ownership of input and output data, and whether secondary use is permitted
- Whether data may be used for training (opt-out)
- Confidentiality, the countries where data is stored, and restrictions on subcontracting
- Notification obligations and allocation of responsibility in the event of a security incident
- Return or deletion of data when the service ends or the contract is terminated

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

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
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

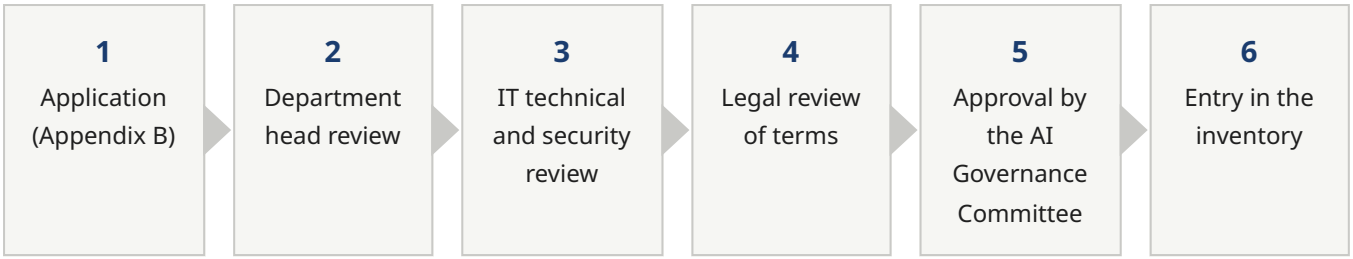

Item	Arrangement
Standard review period	Usually about two weeks from receipt of the application
High-risk uses	Attach the results of the risk assessment (Chapter 13) and the impact assessment. The AI Governance Committee reviews each case individually
If the application is rejected	The applicant is told 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.

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

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 work rules and consideration of legal action

The AI Governance Committee determines whether a violation has occurred. Disciplinary action is decided under the work rules 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 guidelines continue to be revised. The dates in this table are as of the time this document was prepared. Check the latest versions before distribution.

Name	Issued by	Notes
AI Guidelines for Business (Version 1.2)	MIC and METI	March 2026. Covers physical AI and AI agents
Act on the Promotion of Research, Development and Utilization of AI-Related Technologies (AI Promotion Act)	Cabinet Office	Promulgated and fully in force in 2025. Sets out principles and promotion measures
Basic Plan for Artificial Intelligence	Cabinet Office	Approved by the Cabinet in December 2025
Guideline on the Procurement and Utilization of Generative AI for the Evolution and Innovation of Public Administration (Version 2.0, DS-920)	Digital Agency	Revised to Version 2.0 on June 12, 2026
Guidelines for Ensuring the Appropriateness of Research, Development and Utilization of AI-Related Technologies	Cabinet Office	Guidelines issued under the AI Promotion Act
Strategic Report on Management Decision-Making and Governance in the AI Era	AI Governance Association (AIGA)	Governance that enables proactive AI use
EU AI Act (Regulation 2024/1689)	EU	In force since August 2024. Transparency obligations apply from August 2026; high-risk obligations postponed to December 2027 or later
Act on the Protection of Personal Information (APPI)	Personal Information Protection Commission (PPC)	Amendment promulgated on July 17, 2026 (Act No. 56 of 2026). Penalty provisions take effect early, on January 17, 2027
Unfair Competition Prevention Act	METI	Protection of trade secrets. Entering information into AI and the requirement of confidentiality management

Copyright Act	Agency for Cultural Affairs	Article 30-4 and other provisions. General Understanding on AI and Copyright in Japan (March 2024)
Contract Guidelines on Utilization of AI and Data	METI	Practical guidance on contracts for providing and using data
ISO/IEC 42001:2023	ISO/IEC	AI management systems
NIST AI Risk Management Framework 1.0	NIST (US)	2023. Generative AI Profile added in 2024
Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Systems	G7	For organizations developing advanced AI 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 Enterprise or Team editions. With personal editions, what you enter may be used for training, which creates a risk of leaking confidential information.
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 Enterprise edition. Anonymize names and other personal information first.
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.
I want AI to handle customer chats.	Make it clear to customers that they are dealing with AI, and provide a way to escalate to a person. High-risk decisions are made by people.
Can I use generated images in products for sale?	Check the commercial license terms of the tool, and before publishing, confirm that the image does not resemble existing copyrighted works.
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